

Bid-Price Heuristics for Unrestricted Fare Structures in Cargo Revenue Management

Lorenzo Castelli, Raffaele Pesenti and Desirée Rigonat

Abstract In the present work we propose two bid-price based heuristic approaches to tackle a stochastic price-oriented demand of air cargo transportation. We assume fares are non-decreasing over time: the earlier the booking, the cheaper the fare. We consider a single-leg flight without overbooking practices or no-show customers. The proposed framework is suited for air cargo carriers providing a unique product to all its price-oriented customers. The business sustainability relies on a significant reduction in fares that would outperform other benefits, an earlier time of delivery above all. Nevertheless, our modelling framework may be easily extended to other modes of cargo transportation, such as maritime, where a given shipment receives the same service regardless the paid fare, which, in turn, only depends on the time the booking request is made.

Keywords Heuristics · Revenue management · Capacity management · Air cargo · Dynamic programming · Bid-price

1 Introduction

Airlines traditionally prevent high fare paying passengers from buying down into lower fare classes by associating restrictions to each fare level. They make different fares correspond to different products whose characteristics fit the needs

L. Castelli (✉) · D. Rigonat
Dipartimento di Ingegneria e Architettura, Università degli Studi di Trieste,
Via A. Valerio 10, 34127 Trieste, Italy
e-mail: castelli@units.it

D. Rigonat
e-mail: desiree.rigonat@phd.units.it

R. Pesenti
Dipartimento di Management, Università Ca' Foscari di Venezia,
Fondamenta San Giobbe 837, 30121 Cannaregio, VE, Italy
e-mail: pesenti@unive.it

of only one class of customers. This fare policy is reasonable in presence, for each class, of a product-oriented demand, interested to a specific fare product and independent of the availability of cheaper services [25]. The emergence of low-cost carriers shows that this assumption of independent demand segments is becoming more and more unrealistic. Indeed, low-cost carriers offer a single type of product to price-oriented passengers that ignore ticketing restrictions and purchase solely on price [6]. Also, these passengers typically exploit the potentiality of the Internet-based distribution channels to compare the fares of several different airlines.

Precise modelling of customer choice behaviour has been a subject of growing interest in recent years [7, 16]. In fact, the application of pricing algorithms that assume independent demand to a non-segmented market gives rise to the spiral-down effect: customers willing to pay a higher fare but accepting a lower one if available are recorded as lower fare demand when the cheaper product is available. Then, forecasts built on these cheaper product sales underestimate demand for higher fare levels and more low fare products than necessary are made available and, consequently, revenues spiral down [10]. To contrast such an effect, the recent literature proposes pricing policies based on Revenue Management (RM) approaches that segment passengers according to their willingness to pay instead of their compliance to restrictions [27].

In the present work we study the application of these RM techniques to the air cargo industry. Specifically, we propose and analyse the performances of two bid-price based heuristic approaches to tackle a stochastic price-oriented demand of cargo shipments.

The RM approaches are essential tools for cargo shipment as the demand in this industry is in general price-oriented, although some segmentation of the demand by offered product still exists. The same shipment may be charged differently depending on the guaranteed delivery time (express delivery vs. standard delivery). In addition, some airlines, e.g., Lufthansa and American Airlines, offer further optional product features that include boarding priority, pick-up time and location at destination.

The air cargo industry accounts for tens of millions of dollars a year in revenue and, according to the International Air Transport Association (IATA), has stabilised and even shown some weak signs of reprise in certain markets (<http://www.iata.org/pressroom>) after the recession following the 2008 crisis.

Price oriented demand has been a much explored field of study in RM in recent years. In particular, Westermann [27] describes how to integrate revenue management and dynamic pricing concepts based on willingness to pay at airlines with different fare structures. Hopperstad and Belobaba [12] introduce seat inventory control schemas in the single-leg case when demand is not independent from fare class. They forecast the total demand at the lowest fare and repartition it to the different higher fare classes by taking into account the passengers' willingness to pay higher fares. Thus they are in the position to compute the booking limits using traditional algorithms. Fiig et al. [11] address the coexistence of restricted and unrestricted fare structures in markets sharing the same leg(s) on a network. Using

simulations, Cléaz-Savoyen [9] shows that the simultaneous application of the approaches described in Fiig et al. [11] and Hopperstad and Belobaba [12] allows a partial mitigation of the spiral down effect in certain markets. Unfortunately, RM approaches used for passenger flights cannot be directly applied to cargo flights. Indeed, Kasilingam [14], Billings et al. [5] and Slager and Kapteijns [22] point out that the structure of demand and services in the two industries present many differences. For example, each passenger requires just one seat, while each cargo shipment consumes capacity in terms of both weight and volume. Passenger demand presents seasonality patterns while the cargo one is usually more erratic, hence the former is easier to forecast than the latter. The number of passenger customers is usually greater than cargo customers. However, the latter ones make larger bookings, so the behaviour of few of them can significantly influence the prices paid by other customers, a condition which is generally not true for passengers.

In bid-price RM policies, threshold, or “bid”, prices are set for each unit of resource. This kind of price setting was first introduced for airlines’ seat booking by Smith and Penn [23] and Simpson [21]. Since then, they have become widely used due to their conceptual simplicity and easiness of implementation. Talluri and van Ryzin [24] give a comprehensive overview of bid-price techniques pointing out the difficulties arising in determining the right bid. Generally speaking determining the optimal bids may be computational cumbersome, even because they may change dynamically as the flight departure times approach. For this reason, Adelman [1] proposes to compute dynamic bid-prices through a Linear Programming (LP) approximate model. A drawback of this approach is that the number of variables grows exponentially. Bijvank et al. [4] aim at improving robustness towards uncertainty in the demand. To this end, they propose three heuristics that exploit scenario-based stochastic programming methods. Pricing policies for a price-oriented demand are also determined through dynamic programming approaches. As an example, Zhang [28] introduces a dynamic programming decomposition approach and shows that it outperforms static bid-fares one even when bids are frequently recomputed. Popescu et al. [19] use a dynamic programming approach to determine the bid-prices in presence of large shipments. Differently, in presence of small shipments they use a bid-price approach whose bids are obtained by approximating the booking requests with passenger arrival models.

Different authors discuss the consequences of imprecise demand models or incomplete demand data in the air cargo industry. Totamane et al. [26] point out that imprecise demand forecasting causes most cargo airlines to operate at an average ratio between the utilized capacity and the total capacity, the so-called load factor, of 50–70 %. To overcome this limit, they propose a learning algorithm, based on a producer/consumer model, which is able to deliver a 9 % revenue improvement. Luo et al. [17] address the problem of defining overbooking policies that take into account that most booking reservation systems do not keep track of unfulfilled requests. In this framework, they develop an overbooking model that, under appropriate assumptions, they prove providing the optimal

overbooking limits. Amaruchkul et al. [3] study the role of asymmetrical information on customers' demand, operating cost, margin and reservation profit. They investigate under which conditions the maximum combined profit of the involved agents can be obtained in the presence of such an asymmetry. The same authors, in a previous work [2], address uncertainty in package volume, whereas Huang and Hsu [13] and Chew et al. [8] deal with uncertainty in the supplied capacity.

The remainder of this paper is organised as follows. In Sect. 2, we formulate the problem of interest using dynamic programming. In Sect. 3, we introduce two heuristic bid price-based approaches. In Sect. 4, we describe the experiments we carried out to compare the performances of the algorithms and we address the analysis of the results. Finally, in Sect. 5, we draw some conclusions.

2 Dynamic Programming Formulation

In this section, we formulate the problem that a revenue manager faces when he tries to maximise a cargo flight expected revenues as a dynamic programming problem.

The revenue manager works within a single user-class framework, i.e., demand is not segmented or restricted. We assume that there exists a finite set F of M fares $F = \{f_1, \dots, f_M\}$ which are unknown to the customer. The revenue manager has to decide which of the M fares to propose to the customer, and this decision is taken a posteriori, i.e., after having known the size of the shipment, and within the limits set by the following marketing rules that we assume to hold:

2.1 Assumptions

1. A shipment can be accepted only if the flight has residual capacity in terms of volume and weight to accommodate it.
2. A shipment must be charged proportionally to its weight, so that the revenue for a shipment is computed as its weight times the paid fare.
3. The company discourages last minute opportunistic behaviour, hence the succession of fares displayed to the users must be non-decreasing as time approaches flight departures.
4. No price negotiation is allowed.
5. Customers are served one at a time.

We observe that Assumption 2 may sound quite simplistic. However, the results that we present in this work generalize trivially when more complex cost functions are considered.

We also stress the following consequences of the above assumptions. Assumptions 1 and 4 imply that a customer is lost if the displayed fare does not

satisfy his willingness to pay or if there is not enough weight and volume capacity to accommodate his shipment. Also, the revenue manager will always display no more than one fare at a time, since customers are price-oriented and hence always purchase the necessary capacity for their shipments at the lowest available fare that is not higher than their willingness to pay.

Hereafter, we denote by:

- $\{0, 1, \dots, t, t + 1, \dots, T\}$ the set of the time instants at which customers can arrive and be served, 0 is flight reservation opening and T is the closing time;
- $t_k \in \{0, \dots, T\}$ the arrival time of the k customer;
- ϕ_t the probability a customer shows up at time t ;
- f the generic fare;
- C_w and C_v the flight storage capacities in weight and volume, respectively;
- (ω, v) the size of the generic shipment, where ω and v are its weight and its volume, respectively; both are integer values less than or equal to (C_w, C_v) ;
- p_{km} the customer k willingness to pay toward a fare $f_m \in M$, that is the probability that a customer arriving at time t is willing to pay $f_m \omega$ to send a package of size (ω, v) ;
- $q_{\omega v}$ the probability that a shipment has size (ω, v) ;
- $J_m(t, w, v)$ the function that returns the expected optimal revenues from time t on under the hypothesis that the revenue manager displays in t a fare f_m and there are residual capacities w in weight and v in volume.

At each time t , being w and v the residual capacities, the revenue manager faces the following dynamic programming problem:

$$J_m(t, w, v) = [(1 - \phi_t) + \phi_t \sum_{(\omega, v) > (w, v)} q_{\omega v}] J_m(t + 1, w, v) + \phi_t \sum_{(\omega, v) \leq (w, v)} q_{\omega v} \max_{j > m} \{ p_{kj} (f_j \omega + J_j(t + 1, w - \omega, v - v)) + (1 - p_{kj}) J_j(t + 1, w, v) \} \tag{1}$$

with final conditions: $J_m(T, w, v) = J_m(T, 0, v) = J_m(T, w, 0) = 0$ for all $0 \leq w \leq C_w$ and $0 \leq v \leq C_v$.

The first term of Eq. (1) r.h.s. states that the expected revenues from t on coincide with the expected revenues from $t + 1$, when no customer shows up in t , or if the size of the shipment exceeds the available capacity, when a customer shows up in t . Differently, the second term states that, when a customer shows up, the revenue manager must choose the fare to display after having seen the size of the shipment. In this situation, the expected revenues in t are given by the sum of the revenues from the current customer (*stage revenues*) and the expected revenues from $t + 1$ on (*revenues to go*). In presence of a customer and being displayed a fare f_j , two situations may occur:

- with probability p_{kj} the current customer accepts f_j , then the stage revenues are equal to $f_j \omega$ and the next customer finds $(w - \omega)$ and $(v - v)$ as available capacities;

- with probability $(1 - p_{kj})$ the current customer refuses f_j , then stage revenues are 0 and the next customer finds w and v as available capacities.

We conclude this section observing that the difference $J_m(t, w, v) - J_m(t, w - \omega, v - \nu)$ represents the opportunity cost at fare f_m of a shipment of size (ω, ν) making a request at time t when capacities w in weight and v in volume are still available, i.e., it is the expected loss in future revenue from using the capacity now rather than reserving it for future use [25]. Accordingly, if $J_m(t, w, v)$ is differentiable then $\partial J_m(t, w, v)/\partial w$ and $\partial J_m(t, w, v)/\partial v$ are the weight marginal opportunity cost and volume marginal opportunity cost, respectively.

3 Bid-Price Heuristics

Solving problem (1) is, in general, impractical from a computational point of view. For this reason, in this section, we present two heuristics based on threshold values called bid-prices. The rationale behind these heuristics is the following. An optimal policy, solution of (1), accepts a shipment if and only if the revenue it generates is larger or equal to its opportunity cost and bid-prices can be fixed as approximate estimations of the marginal opportunity costs. On the base of this property, bid-price heuristic policies accept a shipment if its revenue is greater or equal to the estimation of its opportunity cost that can be derived from the bid-prices [25]. More formally, we denote by $\pi_w(w, t)$ and $\pi_v(v, t)$ the bid-prices for weight and volume, respectively, when capacities w in weight and v in volume are still available at time t . Then, the opportunity costs of such a request are approximated by $\pi_w(w, t)\omega + \pi_v(v, t)\nu$. Hence a booking request generating revenues r is accepted if and only if [18]:

$$r = f\omega \geq \pi_w(w, t)\omega + \pi_v(v, t)\nu, \quad (2)$$

where f is the fare applied.

3.1 Static BP Heuristic (SBP)

We define bid-prices as static if they are fixed at the beginning of the booking period, i.e., they do not change over time and do not depend on the remaining capacity: $\pi_w(w, t) = \pi_w$ and $\pi_v(v, t) = \pi_v$.

Once the bid-prices are chosen, our heuristic fixes a same fare f for all customers. Specifically, f is set equal to the first accepted shipment's minimum available fare such that the acceptance rule (2) is almost surely respected; that is

$$f = \min\{f_j : f_j\omega_h \geq \pi_w\omega_h + \pi_v\nu_h \text{ and } p_{hj} = 1\}$$

where, h denotes the first customer for which a fare $f_j \in F$ satisfying the above condition exists. The uniqueness of f trivially guarantees that the displayed fares do not decrease over time. However, an evident drawback of this choice is that actual generated revenues depend on the willingness to pay of the first accepted customer k . Heuristic SBP may perform poorly when demand is inverse, especially if bid-price values are low. This performance problem holds true for static bid-price approaches in general since they accept a request as long as its revenues are higher than the computed threshold, without taking into account that the marginal value of the remaining capacity increases over time. A common solution to this is to recalculate bid-prices at pre-defined time intervals (which become smaller as flight booking closure grows closer) in order to take into account the increased marginal value of remaining capacity.

To fix the values of the static bid-prices our heuristic averages the optimal static bid-prices of a set of training instances (possibly based on historical data). Specifically, paralleling the work in Pak and Dekker [18], we combine the findings in Lenstra et al. [15] and in Rinnooy Kan et al. [20], and we observe that, if the demand is known in advance, the choice of accepting shipments of size (ω_k, v_k) generating (potential) revenue $r_k = f^* \omega_k$, where f^* is the maximum fare that customer k is willing to pay, can be made by solving a multi-dimensional knapsack problem. Then, we compute the bid-prices for each training instance by solving the associated knapsack problem through the following greedy algorithm.

Greedy algorithm. For each customer k we define the ratio

$$\delta_k = r_k / (\lambda \omega_k + \mu v_k)$$

where λ and μ are appropriate positive multipliers. The ratio δ_k is a measure of the revenue for unit of the overall resources used by shipment k . The multipliers λ and μ are necessary to express volume and weight capacity requirements as single resource consumption. Items are then ordered by decreasing values of δ_k and accordingly inserted into the knapsack as long as there is available capacity. Clearly, the choice of the multipliers may affect the sequence of shipments entering the knapsack and thus the associated final revenues. However, Rinnooy Kan et al. [20] prove that there exists a pair of multipliers λ^* and μ^* maximising the revenues that can be computed in $O(n^3 \log n)$, where n is the number of shipments.

Let δ^* be the ratio value associated to the last shipment inserted in the knapsack when the multipliers are λ^* and μ^* . We define the instance optimal static bid-prices as:

$$\pi_w = \delta^* \lambda^* \quad \pi_v = \delta^* \mu^* .$$

Indeed, is trivial to observe that, in the above procedure, the shipment of a customer k is inserted into the knapsack if and only if $\delta_k \geq \delta^*$ and there is enough capacity. Under these circumstances, condition (2) is met, as $r_k \geq \delta^* (\lambda^* \omega_k + \mu^* v_k) = \pi_w \omega_k + \pi_v v_k$. □

3.2 Dynamic BP Heuristic (DBP)

In our dynamic bid-price heuristic bid-prices are updated as the requests arrive, in order to capture and exploit the increasing willingness to pay of the demand. Let L_k and M_k be the weight and volume (dynamic) bid-price values, respectively. Hence the acceptance rule at time t_k is $r_k \geq L_k \omega_k + M_k v_k$.

After each accepted request, static bid-prices are calculated by running SBP on the remaining capacity and demand. Let $\underline{\pi}_w(w, t)$ and $\underline{\pi}_v(v, t)$ be the weight and volume bid-prices respectively, when SBP is run on $N - k + 1$ customers with remaining capacities w and v .

Initially, we set: $L_1 = \underline{\pi}_w(C_w, t_1)$, $M_1 = \underline{\pi}_v(C_v, t_1)$. Then we update the dynamic bid-prices according to the following policy:

- $L_{k+1} = L_k$, $M_{k+1} = M_k$: if request k is rejected or if one or both static bid-prices turn out to be lower than or equal to current values of L_k or M_k , that is, either $\underline{\pi}_w(w, t_{k+1}) \leq L_k$ or $\underline{\pi}_v(v, t_{k+1}) \leq M_k$.
- $L_{k+1} = \underline{\pi}_w(w, t_{k+1})$, $M_{k+1} = \underline{\pi}_v(v, t_{k+1})$ if request k is accepted and static bid-prices turn out to be greater than or equal to current values of L_k or M_k , that is $\underline{\pi}_w(w, t_{k+1}) > L_k$ and $\underline{\pi}_v(v, t_{k+1}) > M_k$.

Finally, we choose the fare to display to customer k as $f = \min\{f_j: f_j \omega_k \geq L_k \omega_k + M_k v_k \text{ and } p_{kj} = 1\}$. If no fare satisfies this last condition, customer k is rejected.

The rationale behind the choice of updating dynamic bid-prices only when both SBP bid-prices simultaneously increase is three-fold. First, in this way, at each update, the threshold given by the pair (L_{k+1}, M_{k+1}) is the optimal one for the remaining capacity and demand, according to the SBP algorithm. Second, we cannot update both the bid-prices to lower SBP bid-prices values as otherwise we could not guarantee that the displayed fare does not decrease over time. Third, we cannot update a single bid-price to a higher value when only one SBP bid-price is greater than the current dynamic bid-prices as otherwise we obtain a too selective policy. Indeed, the acceptance threshold increases very rapidly over time since it is updated whenever at least one bid-price augments.

We finally point out that the choice to consider updating dynamic bid-prices only after a request is accepted, and not after any request, is due to reducing computational time.

4 Experimental Results

In this section, we assess the quality of the heuristics introduced in the previous section. The revenues obtained with the two heuristics are then compared also with the ones obtained solving problem (1) under the assumption that the number of customers and their characteristics, in terms of both arrival times and the shipment

sizes, are known in advance. In fact, under this (relaxed although unrealistic) assumption, problem (1) becomes computational tractable and provides an upper bound for the optimal expected revenues $J_I(\theta, C_w, C_v)$. Hereafter, we indicate this last kind of revenues as obtained through dynamic programming (DYM).

4.1 Experiment Design

Each shipment is characterised by three attributes: weight, volume and willingness to pay of its owner; they are expressed in kilos (Kg), cubic meters (m^3) and US dollars (\$), respectively. As it is common practice within the air cargo industry, we distinguish between small and large shipments: small when its weight is between 2 and 45 kg and large when it is between 46 and 500 kg. This distinction is generally applied because of the different weight and volume capacities reserved on the aircraft and the dissimilar fares applied by the carriers. In principle, fares per unit of weight decrease as the weight increases and roughly depend on the distance to be flown. Since in this work all the customers book for the same single-leg flight, we can neglect the dependence on the distance.

Hence, fares are only related to the shipment weight category and to the time of the booking request. We considered four different fares ($f_1 < f_2 < f_3 < f_4$) for both small and large shipments. As anticipated, fares are non-decreasing over time and only one fare is available at each time instant.

In order to deal with uncertainty regarding shipment volume [2, 22], it is usual practice within airlines to associate an average volume for a given weight, i.e., a shipment k of weight ω_k has an average volume $\gamma\omega_k$ where γ is a constant. We set $\gamma = 0.00581$ as in Pak and Dekker [18]. To add variability to the volume, its value v_k is randomly chosen in an interval centred in $\gamma\omega_k$ whose length becomes larger as the weight of the shipment increases.

Tests are run with willingness to pay of the demand either random or inverse. In the former case, the willingness to pay is stationary over the time. Indeed the willingness to pay of customer k is chosen randomly in the interval $[f_1\omega_k - l_m, f_4\omega_k + l_M]$ with $l_m = 0.92$ and $l_M = 2.26$ for small shipments and $l_m = 16.56$ and $l_M = 21.26$ for large shipments. In the latter case (i.e., inverse demand), we introduce a dummy fare $f_0 < f_1$ and we divide the N potential customers in four intervals of approximately equal length. Each customer in interval i , ($i = 1, 2, 3, 4$) randomly chooses between fare f_{i-1} and fare f_i with equal probability $p = 0.5$.

By averaging the values presented by airlines on their websites, we fixed the relevant data as reported in Table 1.

Table 1 Experiment parameters

<i>Small shipments</i>	
Weight capacity C_w	4000 kg
Volume capacity C_v	26.7 m ³
Fares	$f_1 = 4.34\$$ $f_2 = 5.31\$$ $f_3 = 6.04\$$ $f_4 = 7.72\$$
Shipment weight ω_k	$2 \leq \omega_k \leq 45$ kg (integer random)
Shipment volume v_k	v_k is a random number between: [$\gamma\omega_k - 0.0051, \gamma\omega_k + 0.0051$] when $\omega_k \leq 9$ [$\gamma\omega_k - 0.0101, \gamma\omega_k + 0.0101$] when $10 \leq \omega_k \leq 45$
Demand amount N	750
<i>Large shipments</i>	
Weight capacity C_w	15500 kg
Volume capacity C_v	80.25 m ³
Fares	$f_1 = 1.97\$$ $f_2 = 2.67\$$ $f_3 = 3.67\$$ $f_4 = 4.03\$$
Shipment weight ω_k	$46 \leq \omega_k \leq 500$ kg (integer random)
Shipment volume v_k	v_k is a random number in: [$\gamma\omega_k - 0.0201, \gamma\omega_k + 0.0601$] when $46 \leq \omega_k \leq 99$ [$\gamma\omega_k - 0.0151, \gamma\omega_k + 0.171$] when $100 \leq \omega_k \leq 299$ [$\gamma\omega_k - 0.0101, \gamma\omega_k + 0.451$] when $300 \leq \omega_k \leq 500$
Demand amount N	450

4.2 Computational Results

We distinguish four different scenarios in accordance with the four different sets of instances used as input demand. Tables 2 and 3 present the average values obtained over the same set of A instances by the different heuristics for each scenario.

Average execution time refers to the training instances for SBP and DBP. The testing instances simply need to check at each booking request whether the threshold is respected or not, which is a very quick operation.

On the other hand, the DYM algorithm does not require a training phase, so the testing instances are optimally solved through dynamic programming and average execution time refers to these latter.

The results provided shows that, at least for the instances considered, both the heuristics are reasonable. Here, we recall that the DYM results are upper bounds on the optimal ones that a revenue manager could obtain solving (1) without a priori deterministically knowing the demand. Static bid-price heuristic DBP, as expected performs better than SBP with inverse demand. Differently, SBP outperforms DBP when the demand is random. This latter observation is not surprising as, given the non-decreasing assumption, the DBP tries to increase the

Table 2 Results for random demand (average values)

	Small shipments, random demand		
	DYM	SBP	DBP
Revenues	2,458,003	2,283,183	2,082,226
Weight LF	0.999	0.964	0.849
Volume LF	0.869	0.839	0.738
Accepted requests (No.)	178	166	152
Running time (s)	552	<1	624
	Large shipments, random demand		
	DYM	SBP	DBP
Revenues	4,752,162	4,641,544	4,421,279
Weight LF	0.820	0.816	0.770
Volume LF	0.997	0.992	0.935
Accepted requests (num.)	51	49	48
Running time (sec.)	64	<1	2

Table 3 Results for inverse demand (average values)

	Small shipments, inverse demand		
	DYM	SBP	DBP
Revenues	2,712,138	2,053,969	2,121,717
Weight LF	0.999	1.000	1.000
Volume LF	0.870	0.870	0.870
Accepted requests (num.)	171	173	174
Running time (sec.)	594	<1	122
	Large shipments, inverse demand		
	DYM	SBP	DBP
Revenues	5,110,171	4,680,050	4,736,291
Weight LF	0.821	0.823	0.824
Volume LF	0.990	0.999	0.998
Accepted requests (num.)	47	48	49
Running time (sec.)	79	<1	5

revenues becoming more and more selective. However, in the case of random demand, the willingness to pay of the customers does not increase over time and then the DBP may reject too many customers. Indeed, in presence of a stationary demand, it may be reasonable to become less selective when we approach the flight closing time. Unfortunately, the implementation of a similar policy could induce an opportunistic behaviour in the demand and modify its statistics.

5 Conclusions

In the present work we introduced the problems that can arise in RM from an imprecise forecast of customer demand. In this context, we focused on customer's *willingness to pay* as a proven robust measure on which to build capacity management algorithms. Referring to a bi-dimensional capacity scenario (i.e., weight and volume, which is the case of air cargo) and considering demand as deterministic at the time of booking, we first introduced an optimal Dynamic Programming model. Then we proposed two bid-prices based algorithms, one, static and one dynamic, and showed through computational tests their performances in terms of average revenues and computational times.

Future development of this work may address further improvements to the BP policies and comparison to DP based heuristics, another solving approach that has drawn a wide interest in literature.

References

1. Adelman, D.: Dynamic bid prices in revenue management. *Oper. Res.* **55**(4), 647–661 (2007)
2. Amaruchkul, D., Cooper, W., Gupta, D.: Single-leg air-cargo revenue management. Technical report, University of Minnesota (2005)
3. Amaruchkul, D., Cooper, W., Gupta, D.: Air cargo capacity contracts under asymmetric information. Working paper, supply chain and operations research laboratory, University of Minnesota
4. Bijvank, M., Haensel, A., Ecuver, P.L., Marcotte, P.: Time-dependent bid prices for multi-period network revenue management problems. Working paper, (2012)
5. Billings, J., Diener, A., Yuen, B.: Cargo revenue optimisation. *J. Pricing Revenue Manag.* **2**(1), 69–79 (2003)
6. Boyd, E., Kalleisen, R.: The science of revenue management when passengers purchase the lowest available fare. *J. Pricing Revenue Manag.* **3**(2), 171–177 (2004)
7. Bront, J., Méndez-Díaz, I., Vulcano G.: A column generation algorithm for choice-based network revenue management. *Oper. Res.* **57**(3), 769–784 (2007)
8. Chew, E., Huang, H., Johnson, E., Nemhauser, G., Sokol, J., Leong, C.: Short-term booking of air cargo space. *Eur. J. Oper. Res.* **174**, 1979–1990 (2006)
9. Cléaz-Savoien, R.: (2005). Airline revenue management methods for less restricted fare structure. MS thesis, Massachusetts Institute of Technology, Cambridge
10. Cooper, W., Homem-de Mello, T., Kleywegt, A.: Models of the spiral- down effect in revenue management. *Oper. Res.* **54**(5), 968–987 (2006)
11. Fiig, T., Isler, K., Hopperstad, H., Cléaz-Svaoyen, R.: Davn-MR: A unified theory of O&D optimization in a mixed network with restricted and unrestricted fare products. In: AGIFORS Reservation and Yield Management Meeting. Cape Town, South Africa, 15–19 May 2005
12. Hopperstad, C., Belobaba, P.: Alternative RM algorithms for unrestricted fare structures. In: AGIFORS Reservation and Yield Management Meeting. Auckland, New Zealand, 28–31 Mar 2004
13. Huang, K., Hsu, W.: Revenue management for air cargo space with supply uncertainty. *Proc. East. Asia Soc. Transp. Stud.* **5**, 570–580 (2005)
14. Kasilingam, R.: Air cargo revenue management: Characteristics and complexities. *Eur. J. Oper. Res.* **96**, 36–44 (1996)

15. Lenstra, A., Lenstra, J., Rinnooy Kan, A., Wansbeek, T.: Two lines least squares. *Ann. Discret. Math.* **16**, 201–211 (1982)
16. Liu, Q., van Ryzin, G.: On the choice-based linear programming model for network revenue management. *Manufact. Serv. Oper. Manag.* **10**(2), 288–310 (2008)
17. Luo, S., Çakanyıldırım, M., Kasilingam, R.: Two-dimensional cargo overbooking models. *Eur. J. Oper. Res.* **197**(3), 862–883 (2009)
18. Pak, K., Dekker, R.: Cargo revenue management: Bid-prices for a 0-1 multi knapsack problem. Technical report, Erasmus University Rotterdam, The Netherlands (2004)
19. Popescu, A., Barnes, E., Johnson, E., Keskinocak, P.: Bid prices when demand is a mix of individual and batch bookings. *Transp. Sci.* **47** (2), 198–213 (2013)
20. Rinnooy Kan, A., Stougie, L., Vercellis, C.: A class of generalized greedy algorithms for the multi-knapsack problem. *Discret. Appl. Math.* **42**, 279–290 (1993)
21. Simpson, R. W.: Using network flow techniques to find shadow prices for market and seat inventory control. Technical report, MIT Flight Transportation Laboratory Memorandum M89-1, Massachusetts Institute of Technology, Cambridge (1989)
22. Slager, B., Kapteijns, L.: Implementation of cargo management at KLM. *J. Pricing Revenue Manag.* **3**(1), 80–90 (2004)
23. Smith, B.C., Penn, C.W.: Analysis of alternate origin-destination control strategies. In: *Proceedings of AGIFORS 28th Annual Symposium*, pp. 123–144. New Seabury (1988)
24. Talluri, K., van Ryzin, G.: An analysis of bid-price controls for network revenue management. *Manag. Sci.* **44**(1), 1577–1593 (1998)
25. Talluri, K., van Ryzin, G.: *The Theory and Practice of Revenue Management*. Springer, New York (2004)
26. Totamane R., Dasgupta A., Rao S.: Air cargo demand modeling and prediction. *Syst. J. IEEE* forthcoming (2012). doi:[10.1109/JSYST.2012.2218511](https://doi.org/10.1109/JSYST.2012.2218511)
27. Westermann, D.: (Realtime) dynamic pricing in an integrated revenue management and pricing environment: An approach to handling undifferentiated fare structures in low-fare markets. *J. Pricing Revenue Manag.* **4**(4), 395–405 (2006)
28. Zhang, D.: An improved dynamic programming decomposition approach for network revenue management. *Manufact. Serv. Oper. Manag.* **13**(1), 35–52 (2010)