CARGO REVENUE MANAGEMENT: BID-PRICES FOR A 0-1 MULTI KNAPSACK PROBLEM

Kevin Pak and Rommert Dekker

ERIM REPORT SERIES *RESEARCH IN MANAGEMENT*	
ERIM Report Series reference number	ERS-2004-055-LIS
Publication	June 2004
Number of pages	24
Email address corresponding author	kpak@few.eur.nl
Address	Erasmus Research Institute of Management (ERIM) Rotterdam School of Management / Rotterdam School of Economics Erasmus Universiteit Rotterdam P.O. Box 1738 3000 DR Rotterdam, The Netherlands Phone: +31 10 408 1182 Fax: +31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

BIBLIOGRAPHIC DATA AND CLASSIFICATIONS		
Abstract	Revenue management is the practice of selecting those customers that generate the maximum revenue from a fixed and perishable capacity. Cargo revenue management differs from the well-known passenger revenue management problem by the fact that its capacity constraint is 2-dimensional, i.e. weight and volume, and that the weight, volume and profit of each booking request are random and continuous variables. This leads to a multi-dimensional on-line knapsack problem. We show that a bid-price acceptance policy is asymptotically optimal if demand and capacity increase proportionally and the bid-prices are set correctly. We provide a heuristic to set the bid-prices based on a greedy algorithm for the multi-knapsack problem proposed by Rinnooy Kan et al. (1993). A test case shows that these bid-prices perform better than the traditional LP-based bid-prices that do not perform well at all for this problem.	
Library of Congress Classification (LCC)	5001-6182	Business
	5201-5982	Business Science
	HB 143.7	Optimizations Techniques
Journal of Economic Literature (JEL)	M	Business Administration and Business Economics
	M 11 R 4	Production Management Transportation Systems
	C 61	Optimizations Techniques
European Business Schools Library Group (EBSLG)	85 A	Business General
	260 K 240 B	Logistics Information Systems Management
	260 A	Operations management
Gemeenschappelijke Onderwerpsontsluiting (GOO)		
Classification GOO	85.00	Bedrijfskunde, Organisatiekunde: algemeen
	85.34 85.20	Logistiek management Bestuurlijke informatie, informatieverzorging
	85.03	Operations research
Keywords GOO	Bedrijfskunde / Bedrijfseconomie	
	Bedrijfsprocessen, logistiek, management informatiesystemen	
	Goederenvervoer, opbrengsten, online systemen, knapzakprobleem	
Free keywords	Revenue Management, Cargo Transportation, On-Line Knapsack, Multi-Dimensional Knapsack	

Cargo Revenue Management: Bid-Prices for a 0-1 Multi Knapsack Problem

Kevin Pak*

Rommert Dekker†

June, 2004

Abstract

Revenue management is the practice of selecting those customers that generate the maximum revenue from a fixed and perishable capacity. Cargo revenue management differs from the well-known passenger revenue management problem by the fact that its capacity constraint is 2-dimensional, i.e. weight and volume, and that the weight, volume and profit of each booking request are random and continuous variables. This leads to a multi-dimensional on-line knapsack problem. We show that a bid-price acceptance policy is asymptotically optimal if demand and capacity increase proportionally and the bid-prices are set correctly. We provide a heuristic to set the bid-prices based on a greedy algorithm for the multi-knapsack problem proposed by Rinnooy Kan et al. (1993). A test case shows that these bid-prices perform better than the traditional LP-based bid-prices that do not perform well at all for this problem.

Keywords: Revenue Management, Cargo Transportation, On-Line Knapsack, Multi-Dimensional Knapsack

* Erasmus Research Institute of Management and Econometric Institute, Erasmus University Rotterdam, The Netherlands

† Econometric Institute, Erasmus University Rotterdam, The Netherlands

1 Introduction

Revenue management originates from the airline industry and is traditionally aimed at selecting those passengers for a flight that maximize the revenue. Different passengers pay different prices because of their time of booking, place of booking, cancellation options, etc. Because the capacity is fixed and has to be sold before the plane takes off, there is a trade off between accepting a booking request with a certain revenue and waiting for a more profitable booking request that may or may not come. Throughout the years, revenue management has seen many other applications beyond the airline industry, in for example the hotel, car-rental and railway industries.

It is well known that the principles of revenue management can also be applied to cargo transportation. Although a large part of the cargo capacity is usually consumed by shipments that are determined by long-term contracts, also a certain part of the capacity is generally kept available for on-the-spot sales that tend to be more profitable per kg. Booking requests for these spot sales come in during the booking period and have to be accepted or rejected as they come in. It is for these booking requests that we consider constructing a booking control policy. The allocation of the cargo capacity over the long-term contracts and the spot sales is an interesting research topic on itself, but we will not go into this problem in this paper.

The problem that we consider in this paper is that of accepting or rejecting the spot sales as they come in during the booking period. Unlike other revenue management papers, we do not consider a limited number of booking classes for which the number of requests can be estimated. Instead, we let each booking request have its own unique profit, weight and volume as is the case in practice. We formulate the problem as a multi-dimensional on-line knapsack problem. The solution technique that we use is a bid-price acceptance policy that is easy to use in practice. We provide a polynomial time algorithm to obtain values for the bid-prices for which computational results show that they perform better than the LP-based bid-prices often used in passenger revenue management. We also compare the bid-price policy to a dynamic approximation scheme.

1.1 Cargo vs. passenger revenue management

Cargo revenue management differs from passenger revenue management in a number of ways. The most important difference is that two passengers who book in the same price class for a flight can be seen as two identical customers, whereas each cargo shipment is unique. The two passengers both take up one seat, have the same ticket options, generate the same revenue, etc. This generally results in revenue management policies, that determine how many passengers to accept in each price class, or whether or not a specific price class is open for booking at a certain time. Also for cargo transportation there are usually a number of different product types, e.g. mail, fresh products, live animals, secured products and door-to-door service. However, the profit generated by a shipment doesn't just depend on the price, but also on the additional costs associated to the shipment, e.g. special packaging, additional trucking and fuel. This means that the weight, volume and profit per kg of a cargo shipment are random variables. They take on continuous values and differ for all shipments. The uncertainty in passenger revenue management lies in the number of passengers that will arrive for each product type, whereas for cargo revenue management each shipment is a unique product type on itself with properties that are not known before the time of booking.

Further, unlike passenger capacity, the cargo capacity available for the spot sales is also uncertain. This is, because this capacity isn't only dependent on the long-term contracts, but also on the weather conditions, the amount of fuel and more. For a combi-plane, also the number of passengers and the weight and volume of their luggage have to be considered. Further, the actual amount of cargo usually deviates from the initial booking request, especially for the long-term contracts. This means that overbooking is an important aspect of cargo revenue management. In this paper, however, we assume that the cargo capacity available for the spot sales is known. This is in line with other revenue management papers, where uncertainties about the capacity and overbooking are usually left out in order to focus on the booking control policy. We notice, however, that common overbooking techniques can be used in combination with the policies that we discuss in this paper.

Passenger revenue management is usually looked at from a network perspective. That is, a passenger that uses multiple connected flights, should be evaluated according to its overall profit to the flight network and not for each flight individually. This is also the case for cargo shipments and the model that we provide in this paper takes this into account.

1.2 Literature

Although cargo transportation is generally recognized as a natural application for revenue management, it hasn't received a lot off attention in the literature. In fact, Kasilingam (1996) is the only paper that we know that concentrates on cargo revenue management. Kasilingam concentrates on air cargo and discusses the differences between passenger and cargo revenue management for an airline company. Besides the differences discussed in the previous section, he points out that cargo has the possibility to be shipped among different routes as long as it gets to its destination in time. He also indicates that the number of positions when working with containers is limited. This can be another capacity restriction. Our experience with KLM Royal Dutch Airlines, however, tells us that positioning is hardly a problem for air cargo, since the airline company generally builds pallets itself instead of using containers. The revenue management model that Kasilingam describes, is very similar to the booking limit (bucket allocation) models known from passenger revenue management. We don't think this kind of model suits the cargo problem since cargo shipments can't be classified into groups with identical properties as passengers can. We model each booking request to be unique, which leads us to a multi-dimensional on-line knapsack problem.

In a series of articles, Kleywegt and Papastavrou investigate what they call the dynamic stochastic knapsack problem (see: Papastavrou et al. (1996), Kleywegt and Papastavrou (1998) and Kleywegt and Papastavrou (2001)). They mention its application to cargo revenue management. However, their models include only one capacity restriction. They choose a dynamic programming approach to the problem which, although theoretically very interesting, is computationally very demanding. Especially when a second capacity restriction is added to this approach, the state space will become intractable for practical use. Our approach is a static but more efficient and practical one.

The bid-price solution technique that we use in this paper, has been discussed extensively by Talluri and van Ryzin (1998) for the passenger problem. They show that bid-prices are not optimal in general, but are asymptotically optimal when capacity and demand increase proportionally. Their model differs from ours in the sense that they divide the booking requests into classes in such a way that two requests in the same class always have the same capacity requirement. This is not the case in our problem where every shipment has a random weight and volume. Nevertheless, we show that the asymptotic optimality holds in our case as well.

2 Problem formulation

The cargo revenue management problem as we define it in this paper, is such that a number of booking requests come in during the booking period. Each booking request is uniquely defined by its weight, volume, profit and the flights it uses. The number of shipments that the airline company can accept on a flight is determined by the weight and volume capacities of the plane. We assume that the capacities are given and that the actual weight and volume of the shipments do not differ from the booked quantities. The decision to accept or reject a booking request has to be made at the moment the request comes in. This decision depends on the remaining capacities, the expected future demand and the properties of the booking request.

We model the booking requests as a sequence of arrivals over time and measure time in discrete intervals counting backwards, i.e. at time 0 the process ends. Because a shipment can use more than one flight to reach its destination, we include multiple flights in our model. Assume there are m flights that have weight and volume capacities given by the vectors $c_w = (c_{w,1}, c_{w,2}, ..., c_{w,m})^T$ and $c_v = (c_{v,1}, c_{v,2}, ..., c_{v,m})^T$. The capacities can not be negative and are adjusted every time a booking request is accepted. Define $J_t(c_w, c_v)$ as the optimal expected revenue that can be generated with t time units to go and capacities c_w and c_v available. We know that $J_t(c_w, c_v)$ must satisfy:

$J_t(c_w, c_v) \geq 0 \qquad \forall\, c_w, c_v, t$

$J_0(c_w, c_v) = 0 \qquad \forall\, c_w, c_v$

$J_t(0, c_v) = J_t(c_w, 0) = 0 \qquad \forall\, c_w, c_v, t.$

Further, we know that $J_t(c_w, c_v)$ is nondecreasing in t, c_w and c_v, since it is never a disadvantage to have more time or capacity available. For a more exhaustive analysis of the expected reward function, we refer to Papastavrou et al. (1996), who discuss various special cases of the problem.

For each booking request, let r denote the profit of the request. Further define the vector $w = (w_1, w_2, ..., w_m)^T$ such that w_j is equal to the weight of the shipment if the shipment uses flight j ($j = 1, 2, ..., m$) and 0 for all other flights. Likewise define the vector $v = (v_1, v_2, ..., v_m)^T$ to reflect the volume requirements of the shipment. Then, when a booking request comes in with t time units left to go, it should be accepted if and only if:

$$r \geq J_{t-1}(c_w, c_v) - J_{t-1}(c_w - w, c_v - v)\,. \tag{1}$$

The left hand side of equation (1) denotes the direct revenue associated with accepting the cargo shipment, whereas the right hand side gives the estimated opportunity costs of the capacities taken up by the shipment.

The decision rule in equation (1) is similar to the one known for passenger revenue management. Except that in the passenger case, there is only one capacity dimension and the size of a demand is restricted to a limited number of integer values (and in most cases even considered to be equal to one). Lautenbacher and Stidham (1999), Lee and Hersh (1993), Liang (1999) and Subramanian et al. (1999), use a dynamic programming approach to the passenger problem for a single flight. Van Slyke and Young (2000) also consider multiple flights, but acknowledge that the complexity introduced by the increased number of dimensions is tremendous. For the cargo problem, that has two capacity dimensions and continuous demand sizes, a dynamic programming approach becomes computationally intractable. In the next section we discuss a dynamic approximation scheme for the problem and in Section 2.2 we construct a static bid-price policy that is well suited for use in practice.

2.1 Dynamic approximation scheme

In this section we construct a dynamic approximation scheme that can be used for the on-line accept/deny decision. We propose a method to approximate the optimal expected revenue that can be generated for a given set of capacities at a given time, i.e. $J_t(c_w, c_v)$. This way, approximations can be computed for $J_t(c_w, c_v)$ and $J_t(c_w - w, c_v - v)$, whenever a booking request comes in.

If the future demand is known, we can easily compute the optimal future revenue by formulating the problem as an IP problem. Assume that we have a given set of n booking requests that define the problem instance W. Let the profits of the booking requests be given by the vector $r = (r_1, r_2, ..., r_n)^T$. Further, define the matrix $W = [w_{ji}]$, such that w_{ji} is equal to the weight of shipment i if that shipment uses flight j and 0 otherwise. Likewise, define the matrix V for the volume requirements of the shipments. Let $J_t^{\Omega}(c_w, c_v)$ be the optimal future revenue that can be generated for problem instance W with t time units to go and capacities c_w and c_v available, then $J_t^{\Omega}(c_w, c_v)$ can be obtained from the following IP problem:

$$
\begin{aligned}
J_t^{\Omega}(c_w, c_v) = \; & \max_x \; r^T x \\
& Wx \le c_w \\
& Vx \le c_v \\
& x_i \in \{0,1\} \qquad \text{for } i = 1, 2, ..., n,
\end{aligned}
\tag{2}
$$

where $x = (x_1, x_2, ..., x_n)^T$ determines whether a booking request is accepted ($x_i = 1$) or not ($x_i = 0$). This model is known to be NP-hard (see Garey and Johnson (1979)). Standard integer programming solution methods, such as branch-and-bound, can be used to solve small problem instances of the model. For bigger problem instances Rinnooy Kan et al. (1993) provide a polynomial time greedy algorithm that provides an asymptotically optimal solution when demand and capacity increase proportionally. We will discuss this algorithm in detail in the Section 2.3.

If the optimal future revenue is computed for a series of simulated demand instances, i.e. $J_t^{\Omega_1}(c_w, c_v)$, $J_t^{\Omega_2}(c_w, c_v)$, ..., $J_t^{\Omega_k}(c_w, c_v)$, we can obtain an approximation for $J_t(c_w, c_v)$ by taking the average value over the simulated instances. Although not as intractable as dynamic programming, doing this every time a booking request comes in, is computationally still very demanding. In the next section, we construct a static bid-price policy that is easy to use in practice. We will compare the results of the dynamic approximation scheme and the bid-price policy in Section 3.

2.2 Bid-price policy

Bid-price acceptance policies are widely used in passenger revenue management. The idea of a bid-price policy is to determine a value for which a unit of capacity can be sold at a certain point in time. This way, the opportunity costs of a booking request can be approximated by the sum of the bid-prices of the capacities it uses. The booking request is only accepted if its profit exceeds the opportunity costs. A bid-price is determined for every dimension of the capacity, which in our case would mean one weight and one volume bid-price for each flight. Optimally, the bid-price is a function of the remaining time and capacity. In practice, however, the function is usually approximated by a fixed value that is re-evaluated at fixed points in time.

When the bid-prices for the different capacity dimensions are held constant for a longer period of time, the approximation of the opportunity costs of a booking request reduces to a linear combination of the capacity requirements of the request. In order to see this, let $m_w = (m_{w,1}, m_{w,2}, ..., m_{w,m})^{\mathrm{T}}$ and $m_v = (m_{v,1}, m_{v,2}, ..., m_{v,m})^{\mathrm{T}}$ be the bid-prices for the weight and volume capacities of the flights. Then, if a booking request comes in with profit r, and capacity requirements w and v, it is accepted under the bid-price policy if and only if:

$$r \geq m_w^T w + m_v^T v \,, \tag{3}$$

which is the sum of the bid-prices of the capacities it uses multiplied by the size of the booking request. Bid-prices are studied extensively for passenger revenue management by Talluri and van Ryzin (1998). They show that, when the bid-prices are set correctly, a static bid-price policy is asymptotically optimal when the capacities and the demand increase proportionally. We show that this also holds for the cargo problem.

In a probabilistic error analysis, Rinnooy Kan et al. (1993) show that a static bid-price policy is also asymptotically optimal for the 0-1 multi-dimensional knapsack problem. They consider a probabilistic version of the problem as given in model (2) by letting r_i, w_{ij} and v_{ij} (i = 1, 2, ..., n and j = 1, 2, ..., m) be independent identically distributed random variables. Assume that the capacities grow proportionally with the number of items, i.e. $c_w = n\boldsymbol{b}_w$ and $c_v = n\boldsymbol{b}_v$, where $\boldsymbol{b}_w$ and $\boldsymbol{b}_v$ are fixed values. Finally, define z_n as the random variable that denotes the optimum solution value of the problem with n items and $z_n(m_w, m_v)$ as the random variable that denotes the solution value when the items are accepted by the bid-price policy that uses m_w and m_v as the bid-prices. Then under certain conditions concerning the probability distributions, they show that the sequence $\{z_n(m_w, m_v) / z_n\}$ converges to 1 with probability one, if m_w and m_v are chosen correctly. For an intuitive clarification, note that it is the combinatorial aspect of the problem that creates a gap between the optimal and the greedy solutions. In fact, if all items were of the same size, the greedy algorithm would provide the optimal solution. Now, by increasing the number of items and the capacity along with it, the size of each individual item becomes less influential and so does the combinatorial aspect of the problem. Eventually, as the number of items goes to infinity, the combinatorial effect dies out.

We argue that the on-line decision problem that we study in this paper, can be formulated in exactly the way as discussed above, which means that the asymptotic result holds for our problem as well. Without loss of generality we can assume that there is a booking request for every decision period, since we can always consider a booking request to have a profit, weight and volume of value 0. We can now interpret the on-line problem as a 0-1 decision problem for every decision period. When we formulate this as a knapsack problem we let the decision variables in the knapsack correspond to the decision periods in the booking process. This way, all random elements of the problem are modeled in the profit and capacity requirement coefficients, which gives us the model studied by Rinnooy Kan et al. (1993). Thus, also for the cargo revenue management problem, a bid-price policy is asymptotically optimal as demand and capacity increase proportionally and the bid-prices are chosen correctly.

2.3 Obtaining bid-price values

In passenger revenue management, bid-prices are often approximated by the shadow prices of the LP-relaxation of the underlying model that determines the number of passengers to accept in each price class. The cargo problem, however, is a 0-1 decision problem for which the LP-relaxation is a very crude approximation. In order to obtain bid-prices, we make use of a greedy algorithm that Rinnooy Kan et al. (1993) present for the multi-dimensional knapsack problem. The idea of the greedy algorithm is to weigh each capacity dimension, and select those items in the knapsack that have the highest profit per weighted capacity requirement. Assume that $a_w = (a_{w,1}, a_{w,2}, ..., a_{w,m})^T$ and $a_v = (a_{v,1}, a_{v,2}, ..., a_{v,m})^T$ are the non-negative weights for the weight and volume capacities on the flights. As before, let r, w and v be the profit and the weight and volume requirements of a booking request. Then the booking requests that are accepted, are those that have the highest ratio:

$$d = \frac{r}{a_w^T w + a_v^T v} \tag{4}$$

In order to understand how the greedy algorithm works, notice that from (4) we know that we want to select those booking requests for which the values w/r and v/r are small. That is, those requests that do not use a lot of capacity for the profit they provide. For a graphical representation, let each booking request correspond to a point in $\Re_+^{2m}$, given by the coordinates $(w_1/r, w_2/r, ..., w_m/r, v_1/r, v_2/r, ..., v_m/r)$. Then, the greedy algorithm can be regarded as a hyperplane with normal vector $(a_{w,1}, a_{w,2}, ..., a_{w,m}, a_{v,1}, a_{v,2}, ..., a_{v,m})$ that moves upward from the origin and accepts booking request in the order that it encounters them. This is depicted for a single flight in Figure 1. For this 2-dimensional case, the hyperplane is reduced to line with the slope a_w/a_v.

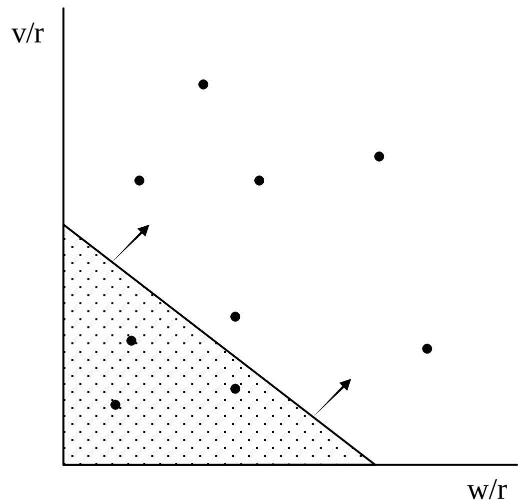

Figure 1. Accepting booking requests for a single flight with the greedy algorithm.

For a given set of booking requests, one can simply select the requests that have the highest d. For the on-line problem, however, where the future booking requests are not known, a threshold value for d has to be specified in advance. If $\hat{d}$ is the threshold value, then a booking request is accepted if:

$$\frac{r}{a_w^T w + a_v^T v} \geq \hat{d}\ , \tag{5}$$

which is the same as:

$$r \geq \hat{d}\,(a_w^T w + a_v^T v)\,. \tag{6}$$

This means that the greedy algorithm can be seen as a bid-price policy, where the bid-prices are set equal to $\hat{d}a_w$ and $\hat{d}a_v$. In terms of Figure 1, a_w and a_v define the slope of the line and the threshold value $\hat{d}$ determines the height of the line.

For a given set of booking requests, Rinnooy Kan et al. (1993) provide a polynomial time algorithm to determine optimal values for the weights and the threshold value when the number of capacity dimensions is not more than the number of booking requests. They show that the problem becomes NP-hard if the number of capacity dimensions exceeds the number of requests. Since there are only two capacity dimensions in the cargo revenue management problem, this is already ruled out as soon as there are two booking requests per flight. Notice that for a given set of booking requests, each set of weights brings forth an ordering in which the requests are considered for acceptance. However, we need only consider those sets of weights that actually induce a different ordering. That is, we only need to consider those sets of weights that construct a sufficiently different hyperplane such that it encounters the booking requests in a different order as it moves up from the origin. In the case of one flight, and thus 2-dimensions, with n booking requests, there are at most $\binom{n}{2}$ of such changes in the ordering.

Lenstra et al. (1982) provide an algorithm that can be used to determine all possible orderings. Looking at Figure 1, the algorithm makes use of the fact that the line that passes through two booking requests, provides exactly that slope for which the requests are swapping places in the ordering. Thus, in order to obtain all orderings, we only have to determine those lines that connect two booking requests. Note that we only have to consider negative slopes, since a positive slope does not provide an ordering of the requests. In fact, whenever a line with a positive slope passes through two booking requests, the request furthest from the origin is always preferred over the other request and no ordering can be based on them changing places in the ordering. Let $x_i = (w/r)_i$ and $y_i = (v/r)_i$ denote the weight and volume per profit for booking request i (i = 1, 2, ..., n). Lenstra et al. (1982) suggest to first order the requests by the sequence $\{x_i\}_{i=1}^{n}$ and then to determine for each i (i = 1, 2, ..., n) and j (j = 1, 2, ..., n) for which $i < j$ (i.e. $x_i = x_j$) and $y_i \geq y_j$, the line that passes through the requests i and j. For each ordering that can be obtained, one can simply compute the solution of the greedy algorithm by accepting the requests until the capacity is full.

Determining the orderings by the algorithm provided by Lenstra et al. (1982) can be done in O(n^2 log n) time. Filling the capacity for a specific ordering costs at most O(n) time. Finally, selecting the ordering that produces the highest profit costs O(log n) time. This means that the total computation time is O(n^3 log n). When we consider m flights, the total number of exchanges is O(n^{2m}). They can be determined in O(n^{2m} log n) time, such that the total computation time is O(n^{2m+1} log n). This is polynomial in n for a given number of flights m. For the case of one flight, a schematic overview of the algorithm is given in Figure 2. This can be easily extended to the case of m flights.

<u>Given</u>: A number of n booking requests defined by $\{(x_i, y_i)\}_{i=1}^{n}$, where x_i and y_i are the weight and volume per profit for booking request i.

<u>Step 1</u>: Order the requests by increasing value of x_i.

<u>Step 2</u>: For $i = 1$ to n:

For $j = i+1$ to n and $y_i = y_j$:

Let $g = \dfrac{y_j - y_i}{x_j - x_i}$.

For $k = 1$ to n: let $h_k = y_k - gx_k$.

Order the requests by increasing value of h_k.

Start accepting requests in this order until no more requests can be accepted.

Let p be the profit obtained and let h be the order value for the last accepted request.

<u>Step 3</u>: Find the maximum profit p^* over all orderings.

Let g^* and h^* be the corresponding slope and order value.

Return: $m_w = -g^*/h^*$ and $m_v = 1/h^*$.

Figure 2. Schematic overview of the algorithm for obtaining bid-prices.

The algorithm discussed above, provides a way to obtain the optimal bid-prices for a given set of booking requests. In practice, however, the bid-prices have to be set before the demand is known. An easy way to obtain bid-prices is to simulate a series of demand sequences, compute the optimal bid-prices for each sequence and set the bid-prices to the average values over all sequences. In the next section, we present some computational results when we apply the bid-price policy and the dynamic approximation scheme to a cargo revenue management test case.

3 Test case

In this section we present computational results for the solution methods described in the previous section when they are applied to a simulated test case. The test case that we consider is a realistic one that reflects the situation that we encountered in practice at KLM Royal Dutch Airlines. We obtain results for the bid-price policy when the bid-prices are determined by: (1) the algorithm discussed in Section 2.3 and (2) the shadow prices of the LP-relaxation of the problem. We compare these results to those of the dynamic approximation scheme and the ex-post optimal solution that can be obtained with hindsight when all demand is known. In the following section we describe the test case. In Section 3.2 we present the computational results.

3.1 Description of the test case

The test case that we construct consists of a single flight with a fixed weight and volume capacity. The booking period is made up of T discrete time intervals of length one. Booking requests come in according to a Poisson arrival process with an arrival rate l. This means that with probability l a booking request is made in a time period and with probability $1-l$ there is not. The total expected number of booking requests is therefore lT and the maximum number of requests is T.

Each booking request has a unique profit, weight and volume, that we denote by r, w and v respectively. The profit, weight and volume are related to each other since they all reflect the size of the shipment. To get around this problem, we define the profit and volume relative to the weight of the shipment and assume the variables w, r/w and v/w to be independent random variables. We assume all booking requests to be independent identically distributed and consider w, r/w and v/w to follow a log-normal distribution. The log-normal distribution generates values that resemble the data that we encountered in practice. We note, however, that we do not try to model the large amount of very small shipments, of for example 1, 2, 10 or 20 kg, that an airline company usually has to deal with. These small shipments are generally highly profitable per kg and take up very little capacity. This means that they are almost always accepted and in practice are often not even subjected to the revenue management decision rule. More than anything, these shipments are used to fill up the wholes in the capacity that are left by the bigger shipments. We choose to exclude the small shipments from the test case such that the algorithm that sets the bid-prices is not influenced by these relatively insignificant shipments. Moreover, note also that it does not seem worthwhile to accurately model shipments of 2 kg next to shipments of 1000 kg. A small error in the latter can easily be greater than the first shipment as a whole. The parameters of the simulation are given in Table 1, whereas a graphical presentation of the probability density functions of w, r/w and v/w up until the 99th percentile are given in Figures 3-5.

Parameter	Value	
T	10000	
l	0.00225	
	Mean	Std. Dev.
w	793.474	942.370
r/w	2.55885	1.39501
v/w	0.00581	0.00338

Table 1. Parameters for the simulation of the demand.

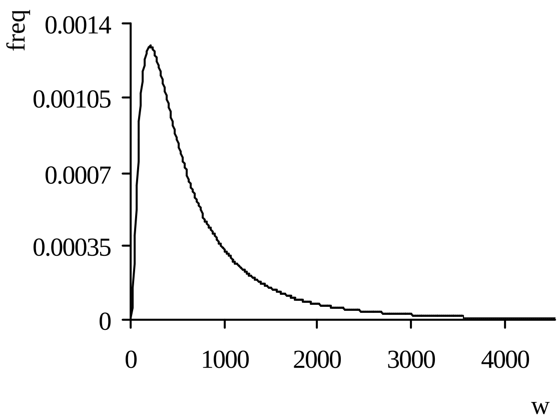

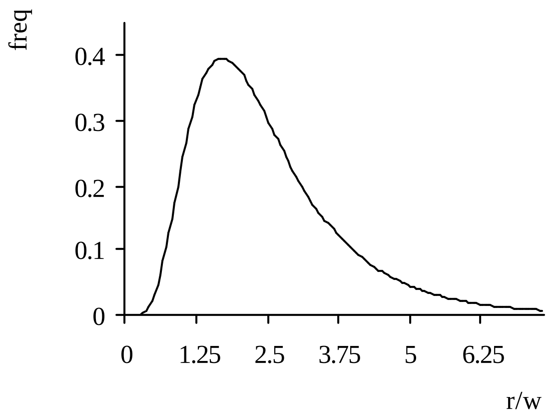

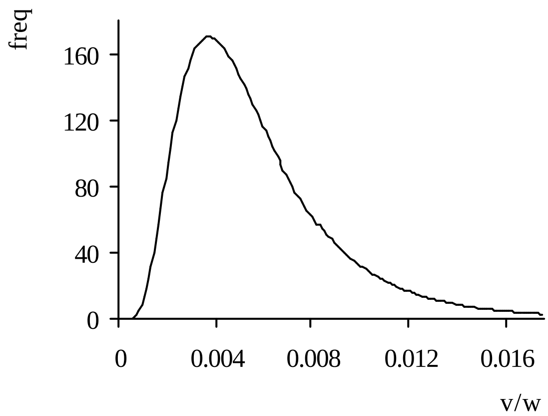

Figures 3-5. Probability density functions for *w*, *r*/*w* and *v*/*w*.

3.2 Computational results

The test case as described above, will be used to compare the performances of the different booking control policies. In order to obtain values for the bid-prices, we simulate 100 demand sequences. For each demand sequence, we obtain the knapsack bid-prices by applying the algorithm discussed in Section 2.3, and LP bid-prices by taking the shadow prices of the LP-relaxation. Taking the average bid-prices over the 100 simulated demand sequences gives us the bid-prices that we will use for the on-line decision problem. This has to be done only once, before the actual booking process starts. The dynamic approximation scheme, on the other hand, needs on-line computation time every time a booking request comes in. In order to limit the computation time, we generate only 10 simulated demand sequences for the dynamic approximation scheme to base its decision on. This has to be done every time a booking request comes in. The dynamic approximation scheme is computationally very demanding and should be seen as an approximation of the optimal on-line booking control policy rather than a practical adversary for the bid-price policies. We also present the ex-post optimal value that could have been obtained with hindsight. This is a natural upper bound for any booking control policy.

In order to compare the performances of the booking control policies, we simulate 100 demand sequences and apply the control policies for a flight that has 10000 kg and 75 m^3 capacity available for the spot sales. These capacities are what you can encounter in practice and correspond to capacity/demand factors of 0.56 and 0.72 for the weight and volume respectively. The computations are performed on a Pentium III 550 MHz personal computer (256 MB RAM), using CPlex 7.1 to optimize the mathematical programming models. An overview of the results is given in Table 2.

	LP Bid-prices	Knapsack Bid-prices	Dynamic Approximation	Ex-post Optimal
Profit				
Average	26491	28725	29748	33555
Std. Dev.	6293	6803	7232	8146
Minimum	15955	16767	15866	19723
Maximum	63136	64004	63810	66579
% of Ex-post Optimal				
Average	80.45	86.58	89.09	100
Std. Dev.	13.00	11.11	8.79	0
Minimum	33.28	45.59	53.41	100
Maximum	100	99.65	100	100
Weight				
Load Factor	0.978	0.951	0.919	0.996
Bid-price	0.190	0.878	-	-
Volume				
Load Factor	0.739	0.651	0.681	0.751
Bid-price	0.868	112.882	-	-
Average On-line Comp. Time (sec.)	0	0	21.86	-

Table 2. Performances of the booking control policies for the simulated test case.

Table 2 shows that the algorithm in Section 2.3 indeed produces better bid-prices than the shadow prices of the LP-relaxation. On average the knapsack bid-prices produce € 2234 more profit than the LP bid-prices. This test case is largely modeled after a daily cargo flight, for which this would add up to more than € 800,000 additional profit per year. Further we see that on average the LP bid-prices obtain 80.45% of the optimal profit that could have been generated with perfect information. The knapsack bid-prices reach up to 86.58% and have a smaller deviation. The dynamic approximation scheme obtains 89.09% with an even smaller deviation. This means that the knapsack bid-prices perform only 2.51% less than the dynamic approximation scheme which we can see as an indication of what an on-line policy without any information on the future can optimally perform. The two bid-price policies have no on-line computation time, whereas the dynamic approximation scheme needs an average of 21.86 seconds to handle a demand sequence. The computation time will grow, however, when we consider more flights or increase the number of simulations that the dynamic approximation scheme uses to estimate the opportunity costs.

Remarkable in Table 2, are the differences between the knapsack and LP bid-prices. The LP bid-prices are much smaller. With hindsight, we can compute an optimal set of bid-prices for each of the 100 simulated demand sequences by applying the algorithm from Section 2.3. This gives us average values of 0.916 and 119.24 for the weight and volume bid-prices. These values are very different from the LP bid-prices of 0.190 and 0.868, but close to the knapsack bid-prices that take on values of 0.878 and 112.88. Note, that the knapsack bid-prices are itself defined as the average values of the optimal bid-prices for a set of 100 simulations. This learns us that the algorithm to obtain values for the knapsack bid-prices is reasonably robust when we use 100 simulations. In Figures 6 and 7 we visualize the LP and knapsack bid-prices next to optimal bid-price values for all simulations. The figures show that the LP bid-prices are situated far below most optimal bid-prices. In fact, it turns out that the LP bid-prices are almost non-restrictive. This means that the policy reduces to nothing more than a first come first serve policy for this test case, which shows the inefficiency of traditional LP bid-prices for cargo revenue management.

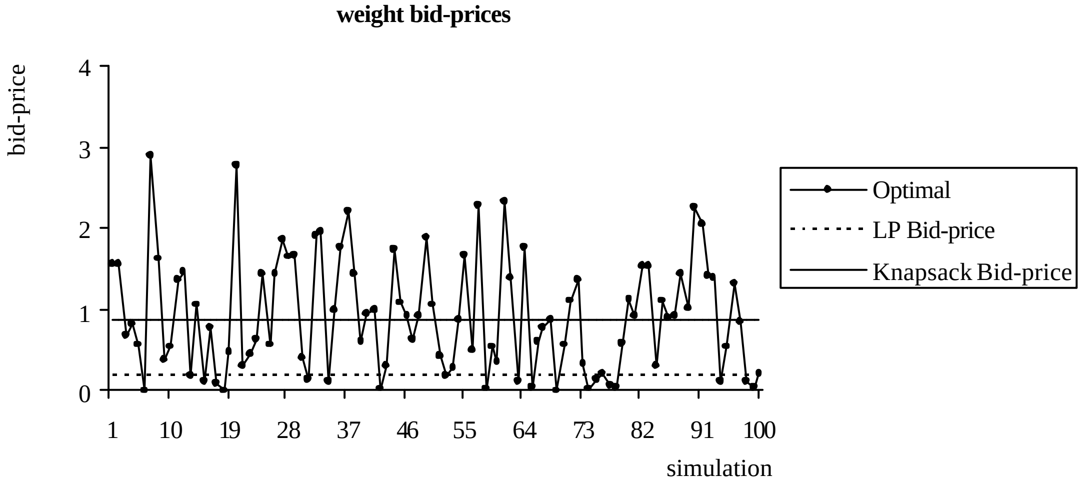

Figure 6: Bid-price values for the weight capacity.

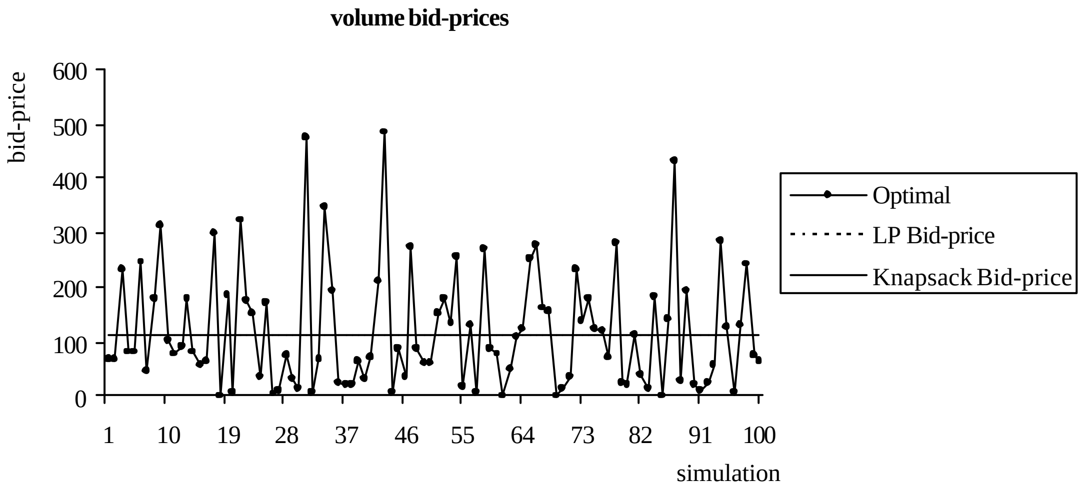

Figure 7: Bid-price values for the volume capacity.

The results discussed above describe one specific test case. It is interesting to see, however, to what extend the performances of the booking control policies are influenced by the set-up of the test case. In the following figures, we show results for some different values for the capacity/demand ratio. We do this by varying the capacity levels while keeping the demand as it is. We define the weight and volume capacities as a percentage of the average demand and shift them from 0.1 to 1.5. Figure 8 shows the profit of the two bid-price policies and the dynamic approximation scheme relative to the ex-post optimal profit for varying values for the weight capacity while keeping the volume capacity fixed at 0.7. Figure 9 shows the same for varying values of the volume capacity while keeping the weight capacity fixed at 0.6.

Each point reflects the average performance of the policy over 100 simulated demand sequences.

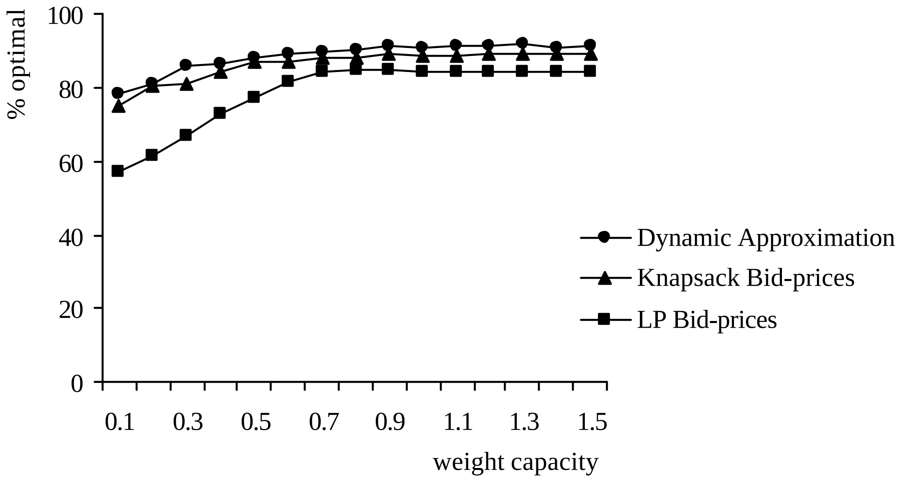

Figures 8. Performances of the booking control policies while varying the weight capacity/demand ratio.

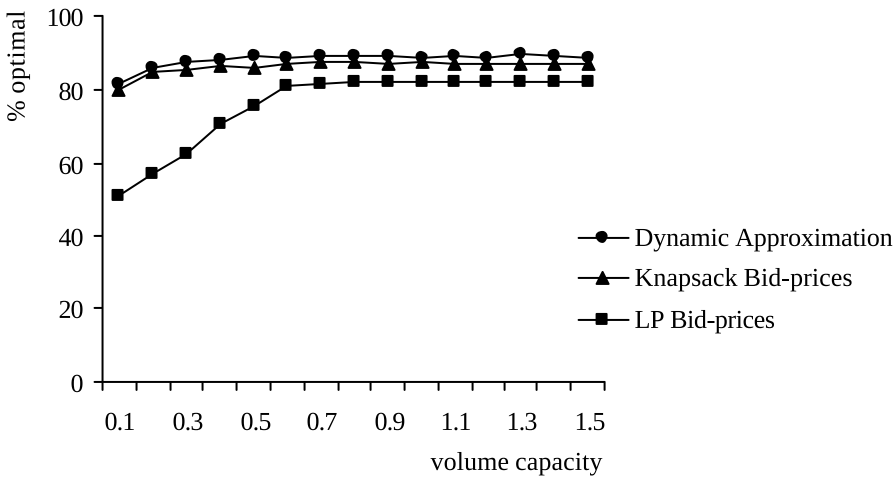

Figures 9. Performances of the booking control policies while varying the volume capacity/demand ratio.

Figures 8 and 9 show that the knapsack bid-price policy outperforms the LP bid-price policy for all cases. The average difference over all capacity combinations between the knapsack bid-price policy and the dynamic approximation scheme is 2% of the optimal profit with perfect information. For the LP bid-price policy this is equal to 11.2%. Especially when one of the capacities is scarce the LP bid-prices tend to perform badly. The performances of the knapsack bid-prices on the other hand follow the performances of the dynamic approximation scheme consistently for all capacity combinations. For a full view of the performances of the booking control policies, we present 3-dimensional plots of the performances when we vary both capacity dimensions at the same time in Figures 10-12.

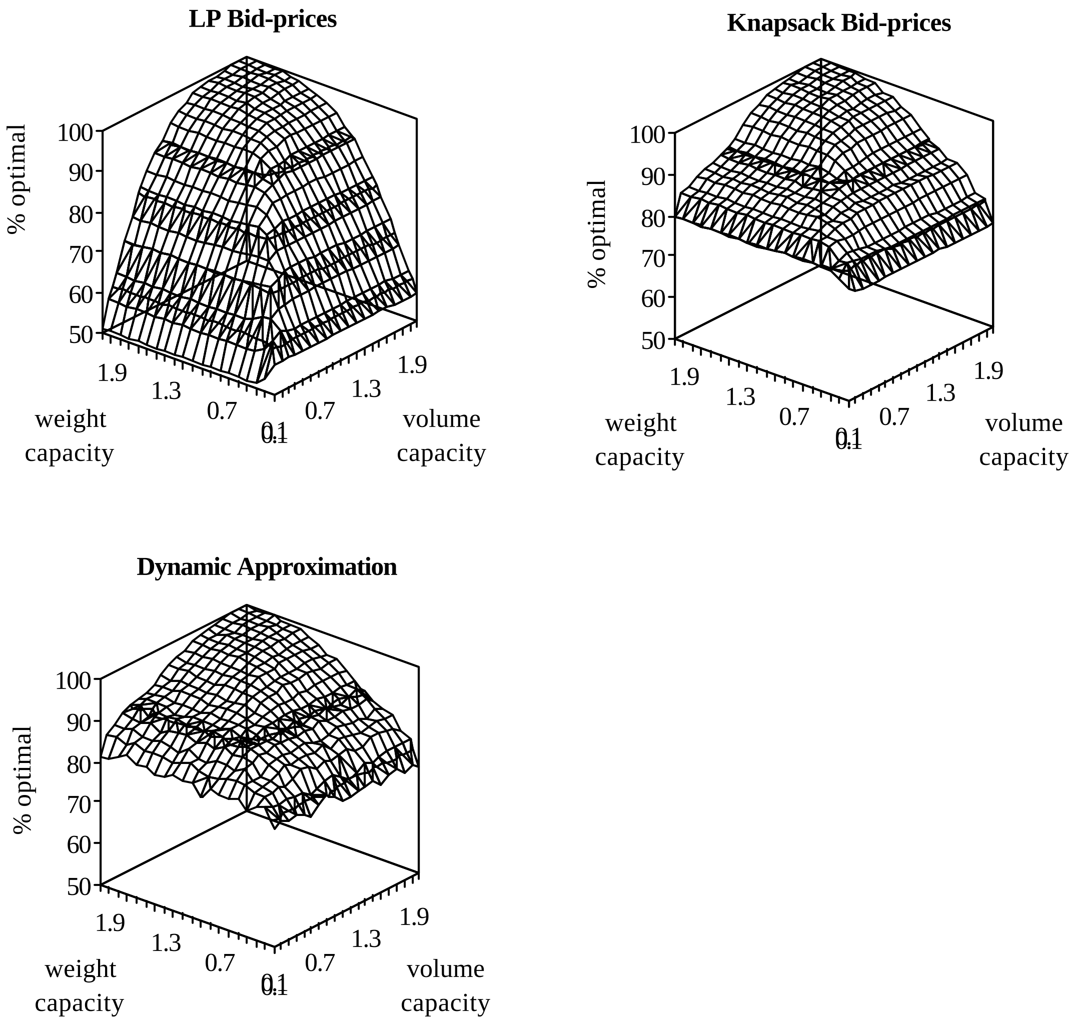

Figures 10-12. Performances of the booking control policies while varying both capacity/demand ratios.

In Figures 10-12 we see that all three booking control policies obtain 100% of the available profit when both capacities are large. This is easy to understand since all booking requests can be accepted in this case. By the steepness of the hillside presented in the figure, we see that the performances of the LP bid-price policy decrease more rapidly than those of the other two policies as the capacities reduce in size and the revenue management decision becomes more important. Note also that the performances of the dynamic policy are more volatile than those of the bid-price policies. The dynamic policy adapts itself to the situation at hand which generally leads to higher profits but also to a more variable kind of decision making. This kind of nervous behavior is another practical drawback of the dynamic approximation scheme. The bid-price policies keep to a fixed policy that is set beforehand and that produces approximately the same results every time it is applied.

4 Conclusion and prospects for future research

Cargo revenue management differs from passenger revenue management in a number of ways. Passengers generally belong to one of a limited number of booking classes and all take up one seat of the total seat capacity. Cargo shipments, on the other hand, are unique in profit, weight and volume. This means that we cannot use traditional revenue management techniques originally constructed for the passenger problem. Instead we formulate the problem as a multi-dimensional on-line knapsack problem. This formulation is capable to include a network of flights.

As for the passenger problem, a bid-price acceptance policy is asymptotically optimal if demand and capacity increase proportionally and the bid-prices are set correctly. We provide a polynomial time algorithm to obtain optimal bid-prices for a given demand sequence. Bid-prices for on-line use can now be constructed by taking the average bid-price values over a number of simulated demand sequences.

A test case based on insights obtained from actual flight data, shows that the knapsack bid-prices outperform the commonly used LP bid-prices for every situation that we consider. Further, the performances of the knapsack bid-prices closely follow the performances of a dynamic approximation scheme that we formulate. Such a dynamic approximation scheme is not convenient in practice, but can give an indication of how well an on-line policy can optimally perform. Bid-prices, on the other hand, are very practical and are already widely used for passenger revenue management.

The cargo revenue management problem as we formulate it in this paper is largely how we encountered it in practice. However, some extensions to the problem that would be worthwhile to examine still remain. First of all, we excluded the overbooking problem from this paper. Overbooking plays an important role in cargo transportation. For cargo, the question is not so much whether or not a booked shipment shows up, but how much the actual weight and volume will deviate from the requested quantities. Overbooking decisions are usually made alongside revenue management decision but it would be interesting to combine the two. Further, as Kasilingam (1996) points out, cargo can be shipped among different routes as long as it arrives at its destination on time. Extending the model to take into account the routing of a cargo shipment, raises a lot of opportunities for the company. Finally, the allocation of the cargo capacity over the long-term contracts and the spot sales would also make an interesting topic for further research. Until now this has largely been a managerial decision. Simulation studies similar to those presented in this paper can be used to estimate the profit that can be generated when a certain amount of capacity is made available for the spot sales. This way, the certainty of a long-term contract can be better evaluated against the additional profit that can possibly be generated by the spot sales.

References

Garey, M.R. and Johnson, D.S. (1979), Computers and Intractability: A Guide to the Theory of NP-Completeness, Freeman, San Francisco, CA.

Kasilingam, R.G. (1996), Air Cargo Revenue Management: Characteristics and Complexities, European Journal of Operational Research, 96, 36-44.

Kleywegt, A.J. and Papastavrou, J.D. (1998), The Dynamic and Stochastic Knapsack Problem, Operations Research, 46, 17-35.

Kleywegt, A.J. and Papastavrou, J.D. (2001), The Dynamic and Stochastic Knapsack Problem with Random Sized Items, Operations Research, 49, 26-41.

Lautenbacher, C.J. and Stidham jr., S. (1999), The Underlying Markov Decision Process in the Single-Leg Airline Yield Management Problem, Transportation Science, 33, 136-146.

Lee, T.C. and Hersh, M. (1993), A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings, Transportation Science, 27, 252-265.

Lenstra, A.K., Lenstra, J.K., Rinnooy Kan, A.H.G. and Wansbeek, T.J. (1982), Two Lines Least Squares, Annals of Discrete Mathematics, 16, 201-211.

Liang, Y. (1999), Solution to the Continuous Time Dynamic Yield Management Model, Transportation Science, 33, 117-123.

Papastavrou, J.D., Rajagopalan, S. and Kleywegt, A.J. (1996), The Dynamic and Stochastic Knapsack Problem with Deadlines, Management Science, 42, 1706-1718.

Rinnooy Kan, A.H.G., Stougie, L. and Vercellis, C. (1993), A Class of Generalized Greedy Algorithms for the Multi-Knapsack Problem, Discrete Applied Mathematics, 42, 279-290.

Subramanian, J., Stidham jr., S. and Lautenbacher, C.J. (1999), Airline Yield Management with Overbooking, Cancellations, and No-Shows, Transportation Science, 33, 147-167.

Talluri, K.T. and van Ryzin, G.J. (1998), An Analysis of Bid-Price Controls for Network Revenue Management, Management Science, 44, 1577-1593.

Van Slyke, R. and Young, Y. (2000), Finite Horizon Stochastic Knapsacks with Applications to Yield Management, Operations Research, 48, 155-172.

Publications in the Report Series Research* in Management

ERIM Research Program: "Business Processes, Logistics and Information Systems"

2004

Smart Pricing: Linking Pricing Decisions with Operational Insights
Moritz Fleischmann, Joseph M. Hall and David F. Pyke
ERS-2004-001-LIS
http://hdl.handle.net/1765/1114

Mobile operators as banks or vice-versa? and: the challenges of Mobile channels for banks
L-F Pau
ERS-2004-015-LIS
http://hdl.handle.net/1765/1163

Simulation-based solution of stochastic mathematical programs with complementarity constraints: Sample-path analysis
S. Ilker Birbil, Gül Gürkan and Ovidiu Listeş
ERS-2004-016-LIS
http://hdl.handle.net/1765/1164

Combining economic and social goals in the design of production systems by using ergonomics standards
Jan Dul, Henk de Vries, Sandra Verschoof, Wietske Eveleens and Albert Feilzer
ERS-2004-020-LIS
http://hdl.handle.net/1765/1200

Factory Gate Pricing: An Analysis of the Dutch Retail Distribution
H.M. le Blanc, F. Cruijssen, H.A. Fleuren, M.B.M. de Koster
ERS-2004-023-LIS
http://hdl.handle.net/1765/1443

A Review Of Design And Control Of Automated Guided Vehicle Systems
Tuan Le-Anh and M.B.M. De Koster
ERS-2004-030-LIS
http://hdl.handle.net/1765/1323

Online Dispatching Rules For Vehicle-Based Internal Transport Systems
Tuan Le-Anh and M.B.M. De Koster
ERS-2004-031-LIS
http://hdl.handle.net/1765/1324

Generalized Fractional Programming With User Interaction
S.I. Birbil, J.B.G. Frenk and S. Zhang
ERS-2004-033-LIS
http://hdl.handle.net/1765/1325

---

* A complete overview of the ERIM Report Series Research in Management:
https://ep.eur.nl/handle/1765/1

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship

Meta-heuristics for dynamic lot sizing: A review and comparison of solution approaches
Raf Jans and Zeger Degraeve
ERS-2004-042-LIS
http://hdl.handle.net/1765/1336

Reinventing Crew Scheduling At Netherlands Railways
Erwin Abbink, Matteo Fischetti, Leo Kroon, Gerrit Timmer And Michiel Vromans
ERS-2004-046-LIS
http://hdl.handle.net/1765/1427

Intense Collaboration In Globally Distributed Teams: Evolving Patterns Of Dependencies And Coordination
Kuldeep Kumar, Paul C. van Fenema and Mary Ann Von Glinow
ERS-2004-052-LIS
http://hdl.handle.net/1765/1446

The Value Of Information In Reverse Logistics
Michael E. Ketzenberg, Erwin van der Laan and Ruud H. Teunter
ERS-2004-053-LIS
http://hdl.handle.net/1765/1447

Cargo Revenue Management: Bid-Prices For A 0-1 Multi Knapsack Problem
Kevin Pak and Rommert Dekker
ERS-2004-055-LIS

A Grounded Theory Analysis Of E-Collaboration Effects For Distributed Project Management
S. Qureshi, M. Liu and D. Vogel
ERS-2004-059-LIS