

Airline Revenue Management with Shifting Capacity

Kevin Pak, Rommert Dekker and Gerard Kindervater

ERIM REPORT SERIES <i>RESEARCH IN MANAGEMENT</i>	
ERIM Report Series reference number	ERS-2003-091-LIS
Publication	2003
Number of pages	25
Email address corresponding author	kpak@few.eur.nl
Address	Erasmus Research Institute of Management (ERIM) Rotterdam School of Management / Faculteit Bedrijfskunde Rotterdam School of Economics / Faculteit Economische Wetenschappen 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	Airline revenue management is the practice of controlling the booking requests such that the planes are filled with the most profitable passengers. In revenue management the capacities of the business and economy class sections of the plane are traditionally considered to be fixed and distinct capacities. In this paper, we give up this notion and instead consider the use of convertible seats. A row of these seats can be converted from business class seats to economy class seats and vice versa. This offers an airline company the possibility to adjust the capacity configuration of the plane to the demand pattern at hand. We show how to incorporate the shifting capacity opportunity into a dynamic, network-based revenue management model. We also extend the model to include cancellations and overbooking. With a small test case we show that incorporating the shifting capacity opportunity into the revenue management decision indeed provides a means to improve revenues.	
Library of Congress Classification (LCC)	5001-6182	Business
	5201-5982	Business Science
	HD 47.3	Cost control
Journal of Economic Literature (JEL)	M	Business Administration and Business Economics
	M 11 R 4	Production Management Transportation Systems
	C 51	Model construction and estimation
European Business Schools Library Group (EBSLG)	85 A	Business General
	260 K 240 B	Logistics Information Systems Management
	255 E	Cost benefit analysis
Gemeenschappelijke Onderwerpsontsluiting (GOO)		
Classification GOO	85.00	Bedrijfskunde, Organiseatiekunde: algemeen
	85.34	Logistiek management
	85.20	Bestuurlijke informatie, informatieverzorging
	85.03	Methoden en technieken, operations research
Keywords GOO	Bedrijfskunde / Bedrijfseconomie	
	Bedrijfsprocessen, logistiek, management informatiesystemen	
	Luchtvaart, rendement, productdifferentiatie, stoelen, dynamische modellen	
Free keywords	Revenue Management, Shifting Capacity, Seat Inventory Control, Dynamic Capacity Management, Convertible Seats	

Airline Revenue Management with Shifting Capacity

Kevin Pak^{*}

Rommert Dekker[†]

Gerard Kindervater[‡]

November, 2003

Abstract

Airline revenue management is the practice of controlling the booking requests such that the planes are filled with the most profitable passengers. In revenue management the capacities of the business and economy class sections of the plane are traditionally considered to be fixed and distinct capacities. In this paper, we give up this notion and instead consider the use of convertible seats. A row of these seats can be converted from business class seats to economy class seats and vice versa. This offers an airline company the possibility to adjust the capacity configuration of the plane to the demand pattern at hand. We show how to incorporate the shifting capacity opportunity into a dynamic, network-based revenue management model. We also extend the model to include cancellations and overbooking. With a small test case we show that incorporating the shifting capacity opportunity into the revenue management decision indeed provides a means to improve revenues.

Keywords: Revenue Management, Shifting Capacity, Seat Inventory Control,
Dynamic Capacity Management, Convertible Seats

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

[†] Econometric Institute, Erasmus University Rotterdam, The Netherlands

[‡] KLM Royal Dutch Airlines, Amsterdam, The Netherlands and Department of Computer Science, Erasmus University Rotterdam, The Netherlands

1 Introduction

Airline companies are confronted with passengers that generate different revenues although they are on the same flight. This comes forth from the fact that the airline companies offer different fare classes and because passengers with different destinations can make use of the same flights. If capacity is scarce, it will therefore be profitable to apply a selection procedure for accepting passengers on the flights. Finding the right mix of passengers on the flights such that revenues are maximized is called revenue management. Revenue management has received a lot of attention throughout the years and has seen applications in a number of industries such as the hotel (see: Baker and Collier (1999), Bitran and Gilbert (1996), Bitran and Monschein (1995), Goldman et al. (2002) and Weatherford (1995)), railroad (see: Ciancimino et al. (1999) and Kraft et al. (2000)) and car rental (see: Geraghty and Johnson (1997)) industries. The main focus of revenue management research, however, remains the airline industry.

In the traditional airline revenue management problem, the capacities of a plane and its different sections, i.e. business and economy class, are fixed. Despite of this, airline companies are not unfamiliar with the practice of shifting capacity from the business to the economy class. This is done by ‘upgrading’ individual passengers from economy to business or by ‘moving the curtain’ between the two sections. A drawback of these procedures is that passengers that pay for the economy class get the luxury of the business class (or at least a business class seat) for free. An airline company should prevent this from happening on a large scale because of the danger that people will anticipate on this and start booking economy class with the probability to be given a business class seat instead of booking business class in the first place. Therefore, upgrading and moving the curtain are not desirable tactics to apply on a large scale.

Another way for shifting business and economy class capacities is provided by so-called convertible seats. By a simple procedure, a row of these seats can be converted from economy class to business class seats and vice versa. When a row is converted from business to economy class, the number of seats in the row is increased and the width of each seat is decreased. The distances between the rows, however, remain the same. An example of this is given in Figure 1. In this figure, taken from

the Swiss Air Lines 2002 Timetable, an Airbus 321 plane is shown equipped with convertible seats. Each row of seats can be used as either six economy class seats or five business class seats. The table included in Figure 1 gives a number of possible configurations of the plane varying from no business class seats in configuration A to 76 business class seats in configuration Q.




Figure 1: Airbus 321 with convertible seats

Because a passenger that has booked for the economy class indeed gets an economy class seat, the drawback previously mentioned for upgrading and moving the curtain does not apply when the convertible seats are used. Moreover, extra seats become available whenever a business class row is converted into an economy class row. The convertible seats can be used without any serious consequences, which makes the plane very flexible in coping with different demand patterns. These different demand patterns can occur among flights that are flown on different days of the week or on different times of the day. Also seasonal differences can be encountered, as can

differences caused by unforeseen events such as a decline in demand during an economic regression.

In this paper, we provide a model to incorporate the shifting capacity opportunity offered by the convertible seats into the traditional revenue management problem. The booking control policy we use to extend for the shifting capacity decision is a dynamic one, which is re-optimized for every new booking request. The underlying model is the standard deterministic programming model for network revenue management. However, as Talluri and van Ryzin (1999), we extend the model to account for the stochastic nature of demand by ways of simulation. Further, we allow for cancellations and overbooking much in the same way as Bertsimas and Popescu (2003) do. Overbooking is not always incorporated in revenue management research, but it is important to do so in combination with the shifting capacity decision. When determining an overbooking policy, one should take into account the fact that one booking can block an entire row of seats from becoming available for the other section of the plane. Also, it is interesting to see if in some cases it is profitable to deny one or two accepted bookings to board on the flight such that the row becomes available for the other section, even though there are costs involved by doing so. For illustration, we describe a test case in which one plane is used for a series of flights and compare the results obtained in a simulated environment when the shifting capacity decision is made (i) beforehand and is kept fixed over all flights, (ii) before each flight, and (iii) dynamically during the booking process of each flight.

The shifting capacity opportunity that we discuss in this paper, is a way to allocate capacity where it is needed. In this respect this paper is the first step towards the integration of revenue management and dynamic capacity management. In the airline industry dynamic capacity management is generally associated with the fleet assignment problem, which is aimed at assigning the different types of planes to the different flights such that revenues are maximized. When this is done dynamically, i.e. when the fleet assignment is changed to match the actual demand when departure time closes, this is also known as demand-driven dispatch (D^3). By using planes equipped with convertible seats, airline companies will be able to fine tune the capacity allocation started with the fleet assignment and will be able to match capacity and demand even better.

2 Problem formulation

The essential decision to be made in airline revenue management is whether or not to accept a booking request when it arrives. In order to make this decision, the direct revenue gained by accepting the request has to be compared to the revenue that can be expected to be gained from the seats if the request is not accepted, i.e. the opportunity costs of the seats. For determining the opportunity costs of the seats, it is important to have a good estimation of the future demand for the various routes and price classes. Further, it should be taken into account that a route requested by a customer can consist of multiple flights. This means that different routes can make use of the same flights. Therefore, in order to get a good approximation of the opportunity costs, the combinatorial effects of the whole network of flights have to be considered.

We formulate the problem under the standard assumptions that the demand is independent over the various routes and price classes and that a rejected booking request is lost forever. The second assumption indicates that the routes and price classes are well differentiated and that a customer will not divert to another route or price class whenever his request is denied. In Section 2.1 we give the traditional formulation for the network revenue management problem. In sections 2.2 and 2.3 we extend the model for the shifting capacity decision and cancellations and overbooking respectively.

2.1 Traditional problem formulation

We assume that the route and price class combinations are well differentiated and can therefore be seen as different products. Then the seat capacities of the flights can be seen as the resources needed for these products. Moreover, when the seat capacities for different sections of the planes are considered to be fixed, the different sections can be considered to be different resources. We model demand as a sequence of booking requests over time and we measure time in discrete intervals counting backwards, i.e. at time 0 the process ends. Define $A = [a_{ij}]$ where $a_{ij} = 1$ if product j uses resource i , and 0 otherwise, for $i = 1, 2, \dots, m$ and $j = 1, 2, \dots, n$. Then, the j^{th}

column of A , a_j , denotes the resources used by one unit of product j . Further, let $c = (c_1, c_2, \dots, c_m)^T$ denote the capacity of each resource, $r = (r_1, r_2, \dots, r_n)^T$ the revenues associated with the products, $u_t = (u_{1,t}, u_{2,t}, \dots, u_{n,t})^T$ the number of products already sold at time t , and $V_t(u_t)$ the optimal expected revenue that can be generated with t time units to go and u_t products already sold. Then, at time t , a booking request for k seats on route and price class combination j will be accepted if and only if:

$$kr_j \geq V_{t-1}(u_{t-1}) - V_{t-1}(u_{t-1} + ka_j). \quad (1)$$

The left hand side of equation (1) denotes the direct revenue associated with the booking request, whereas the right hand side denotes the estimated opportunity costs of the seats taken up by the request.

The decision rule given in (1) has been derived in some form or another for both discrete as well as continuous time by a number of people, including Bertsimas and Popescu (2003), Chen et al. (1998), Lautenbacher and Stidham (1999), Lee and Hersh (1993), Liang (1999), Subramanian et al. (1999), Talluri and Van Ryzin (1998, 1999) and Van Slyke and Young (2000). The difficulty, however, is to approximate $V_t(u_t)$. Let $d_t = (d_{t,1}, d_{t,2}, \dots, d_{t,n})^T$ be the remaining demand with t time units to go. Then if d_t is known, $V_t(u_t)$ can be defined as follows:

$$\begin{aligned} V_t(u_t) = \max_x & r^T x \\ & Ax \leq c \\ & u_t \leq x \leq u_t + d_t \quad \text{integer,} \end{aligned} \quad (2)$$

where $x = (x_1, x_2, \dots, x_n)^T$ determines the number of requests accepted for each route and price class combination. Model (2) provides a linear system of equations that can be solved by standard IP optimization techniques such as branch and bound.

However, d_t will not be known. One way to obtain an approximation for $V_t(u_t)$ is to replace d_t with its expected value. This does, however, not take into account the stochastic nature of demand. A stochastic model has been proposed by Wollmer (1986). But this model is computationally intractable because of the large number of decision variables. A reduced version of Wollmer's model, which considers only a limited number of possible outcomes for the demand, is proposed by De Boer et al.

(2002). The simplest method to incorporate the stochastic nature of demand is probably proposed by Talluri and Van Ryzin (1999). They simulate a sequence of realizations of d_t and compute $V_t(u_t)$ for each of these realizations by applying the model given in (2). Then they approximate $V_t(u_t)$ by averaging the outcomes.

In practice, airline companies do not optimize a model every time a new booking request is made. Instead, they re-optimize the model a number of times during the booking process and use heuristics based on the solution of the model in between optimizations. There are two widely used methods to do this. The first is called the nested booking limit policy. This method uses the solution of the model as the number of seats available for each route and price class combination. Because it is not profitable to reject a request when seats are still available for less desirable route and price class combinations, nested booking limits are defined that allow each route and price class combination to make use of the seats of all the route and price class combinations that are valued lower than itself. The second booking control policy is the bid price policy. For this policy, the opportunity costs of a seat on a flight is approximated by the shadow price of the corresponding capacity constraint. A booking request is only accepted when its revenue is more than the opportunity costs of the seats it uses. These methods are well known both in research as in practice. In this paper we do not use such heuristic policies in between optimizations, but re-optimize the model for every booking request. This gives us the opportunity to focus on the shifting capacity opportunity without any influences of the chosen heuristic.

2.2 Problem formulation with shifting capacity

In this section we extend the standard airline revenue management problem with the shifting capacity decision. We use the same notation as in the previous section with the difference that the capacity c is no longer a constant, but is now dependent on the shifting capacity decision. Assume that each plane has got a limited number of possible capacity configurations collected in the state space Y . Let l be the number of flights and $y = (y_1, y_2, \dots, y_l)^T \in Y$ be the shifting capacity vector which denotes the chosen capacity configuration for each plane. Further, let the capacities be defined as a function of y , $c(y)$. Then, for a given demand vector d_t , $V_t(u_t)$ can be obtained by:

$$\begin{aligned}
V_t(u_t) = \max_{x,y} & r^T x \\
& Ax \leq c(y) \\
& u_t \leq x \leq u_t + d_t \quad \text{integer} \\
& y \in Y,
\end{aligned} \tag{3}$$

where x determines the number of requests accepted for each route and price class combination and y determines the configurations of the planes.

Unless $c(y)$ and Y are of a very specific form, the model provided by (3) will not be a system of linear equations. Therefore, it can be very hard to optimize the model. However, we show that the specifications of $c(y)$ and Y that can be encountered in practice, are such that model (3) reduces to a system of linear equations. In order to see this, we describe the situation of a plane that has got two sections; a business and an economy class section. Then a plane which is equipped with convertible seats usually has a number of seats which are fixed for both sections along with a number of rows which can be used as either business or economy class rows. Assume that the fixed seat capacity for the business class is given by c_b and for the economy class by c_e . Further, let there be R rows of convertible seats which can each be used for either b_b business class seats or b_e economy class seats. Then we can define:

$$c(y) = \begin{pmatrix} c_b + b_b \cdot y \\ c_e + b_e \cdot (R - y) \end{pmatrix}, \quad \text{with } Y = \{y \in \mathbb{N} : 0 \leq y \leq R\}.$$

In this case, y denotes the number of convertible rows appointed to the business class section.

In order to present the model presented in (3) with this specific formulation of $c(y)$, we let c_b , c_e , b_b , b_e and R be vectors of dimension $l \times 1$, such that they contain the shifting capacity information for all flights in the network. Further we also partition r , x , A , u_t and d_t into a part that contains the information concerning the business class and a part that contains the information concerning the economy class. Then we can define:

$$\begin{aligned}
V_t(u_t) = \max_{x_b, x_e, y} & \quad r_b^T x_b + r_e^T x_e & (4) \\
A_b x_b - b_b^T y & \leq c_b \\
A_e x_e - b_e^T (R - y) & \leq c_e \\
u_{b,t} \leq x_b \leq u_{b,t} + d_{b,t} & \quad \text{integer} \\
u_{e,t} \leq x_e \leq u_{e,t} + d_{e,t} & \quad \text{integer} \\
0 \leq y \leq R & \quad \text{integer,}
\end{aligned}$$

where x_b and x_e determine the number of requests accepted for each route and price class combination in the business and economy class and y determines the configurations of the planes. Model (4) consists of a linear system of equations and can therefore be optimized by the same procedures as the model presented in (2). Furthermore, the model has the same number of capacity constraints as the model in (2) and has got only l more decision variables, where l is the number of flights in the flight network. Note that, although model (4) provides a configuration of the planes for every time it is optimized, only at the end of the booking period, at time 0, the planes will be physically converted into the desirable configuration.

2.3 Problem formulation with cancellations and overbooking

In the airline industry, a large amount of bookings typically get cancelled before departure. Therefore, in order to prevent the planes from taking off with empty seats, airline companies overbook the flights. Whenever overbooking is applied, there is a probability that not all bookings can get on the plane. This can happen intentionally when a low fare booking is denied boarding in favor of a high fare booking, or accidentally when the number of cancellations is overestimated. However, there will be a penalty cost involved with denying an accepted booking to board. These can consist of all kinds of costs such as accommodation costs or loss of goodwill. The penalty costs normally prevent airline companies from taking too much risk with overbooking. It is interesting to see, however, if it is worthwhile to take more risk of

bearing the costs of a denied boarding if this means that the entire row becomes available for the other section of the plane.

A deviation from the formulation before, is that we define \bar{x}_t , \bar{u}_t and \bar{d}_t as the net values of x_t , u_t and d_t , where we define the net value as the number of booking requests corrected for the number of cancellations. So if at time t , 30 booking requests have been accepted for route and price class combination j , of which 6 will be cancelled in the future, then $u_t = 30$ but $\bar{u}_t = 24$. Obviously it is not known in advance which bookings will be cancelled. However, we can substitute \bar{u}_t and \bar{d}_t by expected or simulated values. Finally we let q_b and q_e denote the penalty costs of each route and price class combination in the business and economy class, such that we can define:

$$\begin{aligned}
V_t(u_t) = & \max_{\bar{x}_b, \bar{x}_e, y, z_b, z_e} r_b^T \bar{x}_b + r_e^T \bar{x}_e - q_b^T z_b - q_e^T z_e & (5) \\
& A_b(\bar{x}_b - z_b) - b_b^T y \leq c_b \\
& A_e(\bar{x}_e - z_e) - b_e^T (R - y) \leq c_e \\
& \bar{u}_{b,t} \leq \bar{x}_b \leq \bar{u}_{b,t} + \bar{d}_{b,t} & \text{integer} \\
& \bar{u}_{e,t} \leq \bar{x}_e \leq \bar{u}_{e,t} + \bar{d}_{e,t} & \text{integer} \\
& 0 \leq y \leq R & \text{integer} \\
& 0 \leq z_b & \text{integer} \\
& 0 \leq z_e & \text{integer,}
\end{aligned}$$

where z_b and z_e determine which bookings in the business and economy class are denied boarding when the plane takes off.

3 Test Case

In this section we present a test case in order to show how the models described in the previous section can be used and what the added value can be of using a revenue management policy that exploits the shifting capacity opportunity offered by the convertible seats. We do this by comparing the performances of three different revenue management policies. One that does not make use of the fact that the capacity

can be shifted between the business and economy class, a second that does, but only before the start of the booking period, and a third that fully integrates the revenue management and shifting capacity decisions. We test the different revenue management policies by ways of simulation. The test case that we construct is chosen to reflect insights obtained from professionals in the airline industry. This means that the results that we obtain in this specific setting give an indication of what one can expect to find in practice. In Section 3.1 we describe the setting of the test case, after which we present the performances of the different booking control strategies without and with cancellations and overbooking in Sections 3.2 and 3.3 respectively.

3.1 Description of the test case

The test case consists of a single flight-leg that is flown three times by the same plane. Each flight is characterized by its own typical demand pattern. These three flights can be interpreted to be the same flight in different seasons, on different days of the week, or on different times of the day. Specifically, we model one base flight together with one flight that has more business class and less economy class passengers, and one flight that has less business and more economy class passengers. The plane that is used for this test case has a total of 35 rows of passenger seats that can all be used for either five business class seats or six economy class seats. This resembles the Airbus 321 depicted in Table 1 very much. We consider two price classes in the business class and four in the economy class. The prices are given in Table 1. Also given in Table 1 is the average demand of each price class for all three flights. Flight 2 is the base flight. For the first flight, the average business class demand is defined as 30% above the business class demand on the base flight, and the average economy class demand is defined as 30% below the average economy class demand on the base flight. For the third flight this is the other way around.

Class	Class Type	Price (\$)	Average Demand		
			Flight 1	Flight 2	Flight 3
1	Business	400	14.3	11	7.7
2	Business	350	36.4	28	19.6
3	Economy	250	22.4	32	41.6
4	Economy	200	30.8	44	57.2
5	Economy	150	51.1	73	94.9
6	Economy	100	43.4	62	80.6

Table 1: Price classes

Next to the fact that the six price classes differ in price and demand level, they also have a specific booking pattern. The arriving booking requests are modeled by a non-homogeneous Poisson process. This is done by partitioning the booking period into ten smaller time periods each with a constant arrival rate. Demand for the two business classes is assumed to realize at the end of the booking period, whereas demand for the two cheapest price classes is modeled to occur at the beginning of the booking period. Graphical presentations of the arrival patterns of the price classes for the base flight are included in Figure 2. For the other two flights, the booking patterns are the same only with different demand levels. For sake of simplicity, we will not model booking requests for multiple seats, but only consider single seat bookings. Cancellations and overbooking will be incorporated into the test case in Section 3.3.

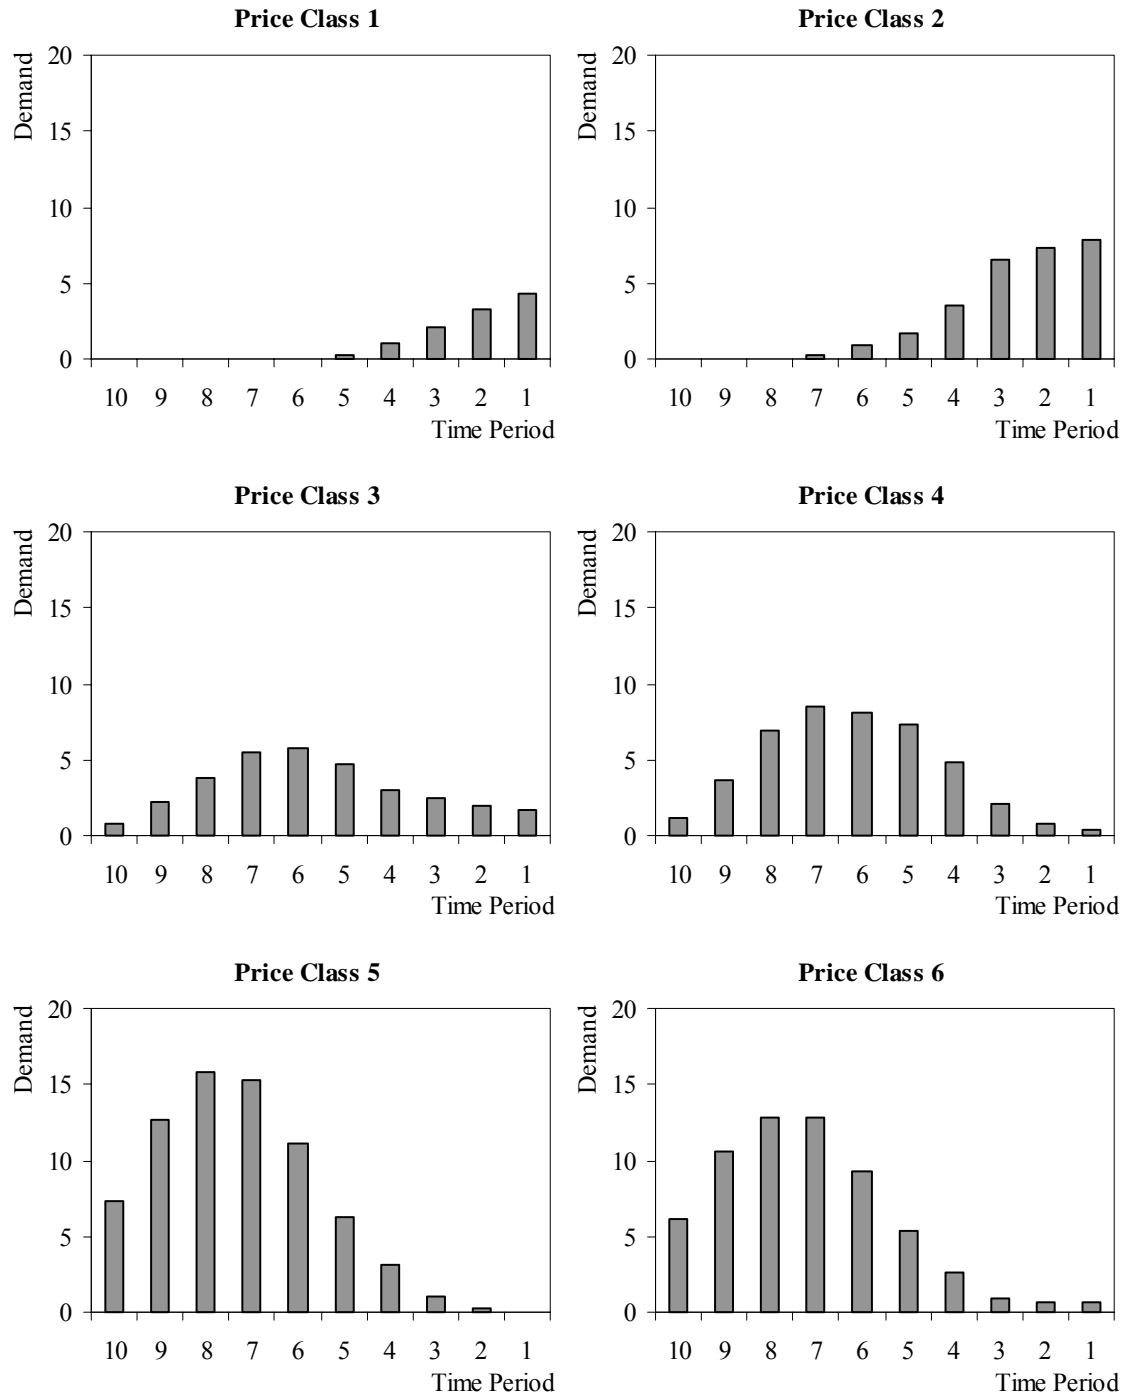


Figure 2: Arrival patterns of the demand for the 6 price classes

3.2 Results without cancellations and overbooking

We compare the performances of three different revenue management policies that differ in the manner in which they deal with the shifting capacity opportunity. All three policies are dynamic policies in the sense that the opportunity costs are estimated anew for every booking request that comes in. The first policy does not incorporate the shifting capacity decision into the revenue management policy. For this policy, the capacity remains fixed over all three flights and the traditional model presented in (2) is used to estimate the opportunity costs. The capacity configuration of the plane is fixed at the configuration obtained when model (4) is used to optimize all three flights at once based on their expected demand. This is done before the booking period, and thus the revenue management process, starts. We will call this policy the Fixed Capacity (FC) policy. The second policy does make use of the shifting capacity opportunity. For each flight, which has its own specific demand pattern, a new capacity configuration is determined for the plane. This configuration is however not changed during the booking period. Before the booking period starts, model (4) is used to determine the capacity configuration based on the expected demand and during the booking period, model (2) is used to estimate the opportunity costs. The third policy makes use of the shifting capacity opportunity dynamically. It fully integrates the shifting capacity and revenue management problems and continually uses model (4) to estimate the opportunity costs. This means that the actual configuration of the plane will be known only at the end of the booking period. We will refer to the second policy as the Shifting Capacity (SC) policy and the third policy as the Dynamic Shifting Capacity (DSC) policy.

All three policies will be applied in both a deterministic and a stochastic manner. This way we obtain six different policies: three deterministic and three stochastic policies. The deterministic policies base their estimation of the opportunity costs on the expected future demand. That is, in the underlying model, either model (2) or model (4), the demand vector is replaced by its expectation. For the stochastic policies ten realizations for the future demand are simulated. For all ten simulations the opportunity costs are determined and the estimation of the opportunity costs that will be used, is defined as the average over the ten cases.

In order to test the booking control policies, we simulate 100 complete booking processes for all three flights. In Table 2 we report the overall performances of the six booking control policies when they are applied to the 100 simulated booking processes. We also give the optimal results that can be determined ex-post for each booking process. The results are generated on a Pentium III 550 MHz personal computer (256 MB RAM), using CPLEX 7.1 to optimize the mathematical programming models. The computation time is measured in seconds.

Policy	Revenue	Standard Deviation	Minimum	Maximum	% Optimal	% Best	Comp. Time
FC_det	121470	3422	113050	129950	94.78	0	3.85
FC_stoch	121688	3479	113150	130100	94.95	1	37.24
SC_det	124957	2873	118550	130200	97.49	5.5	3.75
SC_stoch	125197	3037	118500	131100	97.67	22.5	34.82
DSC_det	125114	2976	118250	130800	97.60	5.5	6.58
DSC_stoch	126032	3048	118500	131700	98.30	65.5	62.50
OPTIMAL	128258	3444	120050	136250	100		

Table 2: Average performances of the booking control policies

In Table 2 we see that the stochastic DSC policy, which is the most sophisticated policy, performs best. On average it reaches up to 98.3% of the optimal revenue that can be obtained and it performs better than all five other policies in 65.5% of the times. The other three policies that make use of the shifting capacity opportunity all obtain revenues that are within 1% of the DSC policy. When the capacity is kept fixed, the revenues that are obtained are clearly less. The extra revenues generated by the shifting capacity opportunity are 2.71% and 2.82% for the deterministic SC and DSC policies respectively, and 2.72% and 3.35% for the stochastic SC and DSC policies respectively. In our small test case this is somewhere between \$1162 and \$1448 per flight. As these flights can be flown one or even multiple times a day, and seeing that the extra revenues can be even more for bigger planes, this can lead to a substantial gain in revenues for the airline company. Table 2 also shows that the differences between the SC and DSC policies are not very large: 0.11% and 0.63% for the deterministic and stochastic policies respectively. This indicates that making the shifting capacity decision before the booking period starts

can be a good alternative to making the shifting capacity decision dynamically. Further we see that treating future demand in a stochastic manner leads to an improvement of the DSC policy of 0.7%. For the FC and SC policies, however, the improvement is much smaller; 0.17% and 0.18% respectively. These improvements seem hardly worthwhile if we consider the additional computation time needed for the stochastic policies. In this case the computation times of the stochastic policies are about 10 times as large as for the deterministic policies, which reflects the fact that we apply the model 10 times to obtain one approximation of the opportunity costs as opposed to one time for the deterministic policies.

In order to see where the differences in the performances come from, we include the average capacity configurations and load factors of the flights for the different booking policies in Table 3. Both for the business and the economy class, this table reports the average number of rows appointed to it, the average number of passengers and the average load factor defined as the number of passengers as a percentage of the total capacity of the plane. We note that for the DSC policy the average number of business and economy rows does not necessarily sum to the total number of rows on the plane because a row which remains empty is not appointed to any of the two sections for this policy. For the FC and SC policies the number of rows appointed to the business and economy classes is known before the booking process starts and is kept fixed no matter whether they are filled or not.

Panel A: Flight 1

Policy	Business			Economy			Total
	Rows	Pass.	Load	Rows	Pass.	Load	Load
FC_det	10	46.59	0.266	25	143.11	0.681	0.948
FC_stoch	10	46.62	0.266	25	143.05	0.681	0.948
SC_det	10	46.59	0.266	25	143.11	0.681	0.948
SC_stoch	10	46.61	0.266	25	143.03	0.681	0.947
DSC_det	11.06	47.86	0.273	23.94	142.06	0.676	0.950
DSC_stoch	11.17	48.57	0.278	23.83	141.46	0.674	0.951
OPTIMAL	11.15	49.42	0.282	23.85	141.96	0.676	0.958

Panel B: Flight 2

Policy	Business			Economy			Total
	Rows	Pass.	Load	Rows	Pass.	Load	Load
FC_det	10	39.70	0.227	25	149.10	0.710	0.937
FC_stoch	10	39.69	0.227	25	148.73	0.708	0.935
SC_det	8	37.65	0.215	27	161.09	0.767	0.982
SC_stoch	8	37.69	0.215	27	160.52	0.764	0.980
DSC_det	7.97	37.18	0.212	27.21	162.00	0.771	0.984
DSC_stoch	8.01	38.23	0.218	26.99	160.92	0.766	0.985
OPTIMAL	8.13	39.55	0.226	26.87	161.22	0.768	0.994

Panel C: Flight 3

Policy	Business			Economy			Total
	Rows	Pass.	Load	Rows	Pass.	Load	Load
FC_det	10	28.71	0.164	25	149.59	0.712	0.876
FC_stoch	10	28.71	0.164	25	149.41	0.711	0.876
SC_det	5	24.41	0.139	30	179.43	0.854	0.994
SC_stoch	5	24.43	0.140	30	179.30	0.854	0.993
DSC_det	5.24	25.17	0.144	29.76	177.42	0.845	0.989
DSC_stoch	5.56	26.79	0.153	29.44	175.62	0.836	0.989
OPTIMAL	5.76	28.15	0.161	29.24	175.44	0.835	0.996

Panel D: Total

Policy	Business			Economy			Total
	Rows	Pass.	Load	Rows	Pass.	Load	Load
FC_det	10	38.33	0.219	25	147.27	0.701	0.920
FC_stoch	10	38.34	0.219	25	147.06	0.700	0.919
SC_det	7.67	36.22	0.207	27.33	161.21	0.768	0.975
SC_stoch	7.67	36.24	0.207	27.33	160.95	0.766	0.974
DSC_det	8.03	36.73	0.210	26.97	160.49	0.764	0.974
DSC_stoch	8.25	37.86	0.216	26.75	159.33	0.759	0.975
OPTIMAL	8.35	39.04	0.223	26.65	159.54	0.760	0.983

Table 3: Average capacity configurations, number of passengers and load factors

Table 3 shows that when the shifting capacity opportunity is exploited, the average number of business class rows ranges from 5 on the third flight with a low business class demand to 10 on the first flight with a high business class demand. The FC policy fixes the number of business class rows at 10 and economy class rows at 25, which is a good configuration when the business class demand is high, but results in empty business class seats when the business class demand is low. This is reflected in the average load factors of the flights. For the first flight, the load factors of the FC policies are the same as those for the SC policies and only a little less than those of the DSC policies. For the third flight, however, the total load factors of the FC policies are more than 11% under those for the SC and DSC policies. Combined over the three flights the load factors of the FC policies are at least 5% less than for the SC and DSC policies.

3.3 Results with cancellations and overbooking

In this section we extend the test case to include cancellations and overbooking. For this we model each booking request to have a probability that it will be cancelled. This cancellation probability is dependent on the price class of the booking request and is the overall probability that the request is cancelled at some time during the time of booking and the end of the booking period. We assume the cancellation probability to be homogeneous over time, such that a booking request that is made t time periods before the end of the booking period and that has cancellation probability p , has a cancellation probability per time unit of p/t . This way we are able to model the cancellations by a homogeneous Poisson process. We model the two business classes to have a cancellation probability of 10%, the first two economy classes of 12.5% and the two cheapest price classes of 15%. The simulated demand is increased proportionally to these percentages in order to keep the net demand on the same level as in the previous section.

As discussed in Section 2.3, the penalty costs of denying an accepted booking to board have to be taken into account when overbooking is allowed. We set the penalty costs at \$500 for all price classes. This is more than the maximum revenue that can be obtained from any price class, which means that it is never profitable to

accept an extra high price booking if this means that another booking has to be denied boarding. With the shifting capacity opportunity, however, it can be profitable to deny one or two bookings to board if this makes the entire row available for the other section of the plane. We simulate 100 booking processes on which we apply the same six booking control policies as in Section 3.2. The overall performances of the six policies are reported in Table 4 together with the optimal results that can be determined ex-post.

Policy	Revenue	Standard Deviation	Minimum	Maximum	% Optimal	% Best	Comp. Time
FC_det	120656	3607	110150	129600	94.20	0	5.65
FC_stoch	120829	3669	110050	129750	94.36	4	52.05
SC_det	123663	2850	116050	128700	96.56	5.5	5.33
SC_stoch	123848	2844	115800	129300	96.70	14.5	49.11
DSC_det	124260	2805	117150	130600	97.02	12	9.86
DSC_stoch	125154	2801	117600	130000	97.71	64	84.41
OPTIMAL	128172	3630	119900	135850	100		

Table 4: Average performances of the booking control policies with cancellations and overbooking

The results presented in Table 4, show that the booking control policies perform less when we consider cancellations and overbooking. This is not because less revenue is available, which is contradicted by the very small differences between the optimal revenues of both cases, but because the cancellations make the problem more difficult. The differences of the performances of the policies with and without cancellations and overbooking can mount up to nearly 1% of the optimal revenue. Apart from this, the differences between the performances of the various policies show very much the same patterns as without cancellations and overbooking. The stochastic DSC policy performs best and the other policies that make use of the shifting capacity opportunity do not stay behind far. The deterministic FC policy performs 2.36% and 2.82% less than the deterministic SC and DSC policies respectively. And the stochastic FC policy performs 2.34% and 3.35% less than the stochastic SC and DSC policies respectively. This means that the extra revenue that

can be obtained by exploiting the shifting capacity opportunity does not change much when cancellations and overbooking are taken into account.

The average number of denied boardings per flight are presented in Table 5. With an average number of denied boardings of 0.537, as opposed to 0.187 to 0.32 for the other policies, the stochastic DSC policy is clearly the least careful policy with regard to overbooking. This is conform our idea that the shifting capacity opportunity can make it profitable in some cases to bear the costs of a denied boarding if this means that the row of seats can be used for the other section of the plane. Note however, that in spite of the penalty costs it endures, the DSC policy is still the most profitable policy with the highest load factors.

Policy	Flight 1	Flight 2	Flight 3	Total
FC_det	0.300	0.320	0.220	0.280
FC_stoch	0.350	0.250	0.150	0.250
SC_det	0.300	0.440	0.220	0.320
SC_stoch	0.330	0.260	0.310	0.300
DSC_det	0.180	0.210	0.170	0.187
DSC_stoch	0.540	0.520	0.550	0.537
OPTIMAL	0	0	0	0

Table 5: Average number of denied boardings per flight

Finally, in Table 6 we present the average capacity configurations and load factors of the flights for the different booking control policies. The capacity configurations and load factors of the flights show no large deviations from those obtained for the case without cancellations and overbooking, except for the fact that the DSC policies tend to appoint some more seats for the business class. This can be seen most clearly for the first flight, where the DSC policies appoint more than a complete row extra to the business class.

Panel A: Flight 1

Policy	Business			Economy			Total
	Rows	Pass.	Load	Rows	Pass.	Load	Load
FC_det	10	47.48	0.271	25	142.41	0.678	0.949
FC_stoch	10	47.61	0.272	25	142.25	0.677	0.949
SC_det	10	47.48	0.271	25	142.41	0.678	0.949
SC_stoch	10	47.60	0.272	25	142.24	0.677	0.949
DSC_det	11.03	48.91	0.279	23.97	142.08	0.677	0.956
DSC_stoch	11.13	49.76	0.284	23.87	141.77	0.675	0.959
OPTIMAL	11.38	50.89	0.291	23.62	140.52	0.669	0.960

Panel B: Flight 2

Policy	Business			Economy			Total
	Rows	Pass.	Load	Rows	Pass.	Load	Load
FC_det	10	39.97	0.228	25	148.82	0.709	0.937
FC_stoch	10	40.00	0.229	25	148.25	0.706	0.935
SC_det	8	37.40	0.214	27	160.88	0.766	0.980
SC_stoch	8	37.44	0.214	27	159.93	0.762	0.976
DSC_det	7.88	37.03	0.212	27.12	161.41	0.769	0.980
DSC_stoch	8.04	37.96	0.217	26.96	160.94	0.766	0.983
OPTIMAL	8.22	39.95	0.228	26.78	160.68	0.765	0.993

Panel C: Flight 3

Policy	Business			Economy			Total
	Rows	Pass.	Load	Rows	Pass.	Load	Load
FC_det	10	28.85	0.165	25	148.91	0.709	0.874
FC_stoch	10	28.85	0.165	25	148.52	0.707	0.872
SC_det	5	23.95	0.137	30	178.62	0.851	0.987
SC_stoch	5	24.18	0.138	30	178.24	0.849	0.987
DSC_det	5.30	24.94	0.143	29.70	176.83	0.842	0.985
DSC_stoch	5.51	26.60	0.152	29.49	175.99	0.838	0.990
OPTIMAL	5.82	28.39	0.162	29.18	175.08	0.834	0.996

Panel D: Total

Policy	Business			Economy			Total
	Rows	Pass.	Load	Rows	Pass.	Load	Load
FC_det	10	38.77	0.222	25	146.72	0.699	0.920
FC_stoch	10	38.82	0.222	25	146.34	0.697	0.919
SC_det	7.67	36.28	0.207	27.33	160.64	0.765	0.972
SC_stoch	7.67	36.41	0.208	27.33	160.14	0.763	0.971
DSC_det	8.07	36.96	0.211	26.93	160.11	0.762	0.974
DSC_stoch	8.23	38.11	0.218	26.77	159.57	0.760	0.978
OPTIMAL	8.47	39.74	0.227	26.53	158.76	0.756	0.983

Table 6: Average capacity configurations, number of passengers and load factors with cancellations and overbooking

4 Conclusion and prospects for future research

In this paper we introduced convertible seats into the airline revenue management problem. These seats create the opportunity to shift capacity between the business and economy class sections of a plane. We formulated a mathematical programming model to account for the shifting capacity opportunity which can be used both in a deterministic and in a stochastic manner. This model is not much harder than traditional network revenue management models and is also extended to incorporate cancellations and overbooking.

We constructed a test case where a single plane is used for multiple flights with different demand patterns. The test case shows that the shifting capacity opportunity gives a rise in revenues of more than 3.3% of the optimal revenue that can be obtained. When the shifting capacity decision is made only once before each flight, the extra revenues are still more than 2.8% of the optimal revenue. The shifting capacity opportunity also increases the load factor of the plane from 92% to more than 97%. When cancellations and overbooking are taken into account these results remain the same. We also observe that taking the shifting capacity opportunity into account can result in a policy that is less careful with respect to overbooking. This is because the opportunity costs of a booking can become very large whenever the booking is

blocking an entire row from becoming available for the other section of the plane. Therefore, in some cases it can be worthwhile to take the risk of a denied boarding. Further we see that a stochastic use of the model increases the performance of the booking control policy over a deterministic use of the model, but never more than 0.7% of the optimal revenue that can be obtained. The computation time of a stochastic policy will however be considerably larger than for a deterministic policy.

This paper provides a way to model the shifting capacity decision and an indication of the added value of doing so. The booking control policies that we construct in this paper are computationally very cumbersome and will not always be applicable in practice in this exact way. Therefore, a study on computationally less demanding booking control policies could prove useful. For this, one can think of bid-prices that serve as approximations of the opportunity costs for a longer period of time or nested booking limits that determine the number of booking requests to accept for each price class. Both can be based on the models introduced in this paper. Further, we acknowledge that our test case is but an initial one and many more can be constructed to obtain further insights. For example, our test case consists of a single flight, which gives us the opportunity to illustrate things more clearly, and does not include multiple seat booking requests. Finally, we would like to mention that most extensions to the standard airline network revenue management problem that are suggested throughout the literature can be applied to the model that we provide in this paper as well. This comes forth from the fact that our model still resembles the standard models very much.

References

Baker, T.K. and Collier, D.A. (1999), A Comparative Revenue Analysis of Hotel Yield Management Heuristics, *Decision Sciences*, 30, 239-263.

Bertsimas, D. and Popescu, I. (2003), Revenue Management in a Dynamic Network Environment, *Transportation Science*, 37, 257-277.

Bitran, G.R. and Gilbert, S.M. (1996), Managing Hotel Reservations with Uncertain Arrivals, *Operations Research*, 44, 35-49.

Bitran, G.R. and Mondschein, S. (1995), An Application of Yield Management to the Hotel Industry Considering Multiple Day Stays, *Operations Research*, 43, 427-443.

Chen, V.C.P., Günther, D. and Johnson, E.L. (1998), A Markov Decision Problem Based Approach to the Airline YM Problem, Working Paper, The Logistics Institute, Georgia Institute of Technology, Atlanta, GA.

Ciancimino, A., Inzerillo, G., Lucidi, S. and Palagi, L. (1999), A Mathematical Programming Approach for the Solution of the Railway Yield Management Problem, *Transportation Science*, 33, 168-181.

De Boer, S.V., Freling, R. and Piersma, N. (2002), Mathematical Programming for Network Revenue Management Revisited, *European Journal Of Operational Research*, 137, 72-92.

Geraghty, M.K. and Johnson, E. (1997), Revenue Management Saves National Car Rental, *Interfaces*, 27, 107-127.

Goldman, P., Freling, R., Pak, K. and Piersma, N. (2002), Models and Techniques for Hotel Revenue Management using a Rolling Horizon, *Journal of Revenue and Pricing Management*, 1, 207-219.

Kraft, E.R., Srikar, B.N. and Philips, R.L. (2000), Revenue Management in Railroad Applications, *Transportation Quarterly*, 54, 157-176.

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.

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

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.

Talluri, K.T. and van Ryzin, G.J. (1999), A Randomized Linear Programming Method for Computing Network Bid Prices, *Transportation Science*, 33, 207-216.

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

Weatherford, L.R. (1995), Length of Stay Heuristics: Do They Really Make A Difference, *Cornell Hotel and Restaurant Administration Quarterly*, 36, 70-79.

Wollmer, R.D. (1986), A Hub-Spoke Seat Management Model, unpublished company report, Douglas Aircraft Company, McDonnell Douglas Corporation, Long Beach, CA.

Publications in the Report Series Research* in Management

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

2003

Project Selection Directed By Intellectual Capital Scorecards

Hennie Daniels and Bram de Jonge

ERS-2003-001-LIS

<http://hdl.handle.net/1765/265>

Combining expert knowledge and databases for risk management

Hennie Daniels and Han van Dissel

ERS-2003-002-LIS

<http://hdl.handle.net/1765/266>

Recursive Approximation of the High Dimensional max Function

Ş. İl. Birbil, S.-C. Fang, J.B.G. Frenk and S. Zhang

ERS-2003-003-LIS

<http://hdl.handle.net/1765/267>

Auctioning Bulk Mobile Messages

S.Meij, L-F.Pau, E.van Heck

ERS-2003-006-LIS

<http://hdl.handle.net/1765/274>

Induction of Ordinal Decision Trees: An MCDA Approach

Jan C. Bioch, Viara Popova

ERS-2003-008-LIS

<http://hdl.handle.net/1765/271>

A New Dantzig-Wolfe Reformulation And Branch-And-Price Algorithm For The Capacitated Lot Sizing Problem With Set Up Times

Zeger Degraeve, Raf Jans

ERS-2003-010-LIS

<http://hdl.handle.net/1765/275>

Reverse Logistics – a review of case studies

Marisa P. de Brito, Rommert Dekker, Simme D.P. Flapper

ERS-2003-012-LIS

<http://hdl.handle.net/1765/277>

Product Return Handling: decision-making and quantitative support

Marisa P. de Brito, M. (René) B. M. de Koster

ERS-2003-013-LIS

<http://hdl.handle.net/1765/278>

* A complete overview of the ERIM Report Series Research in Management:
<http://www.irim.eur.nl>

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

Managing Product Returns: The Role of Forecasting
Beril Toktay, Erwin A. van der Laan, Marisa P. de Brito
ERS-2003-023-LIS
<http://hdl.handle.net/1765/316>

Improved Lower Bounds For The Capacitated Lot Sizing Problem With Set Up Times
Zeger Degraeve, Raf Jans
ERS-2003-026-LIS
<http://hdl.handle.net/1765/326>

In Chains? Automotive Suppliers and Their Product Development Activities
Fredrik von Corswant, Finn Wynstra, Martin Wetzels
ERS-2003-027-LIS
<http://hdl.handle.net/1765/363>

Mathematical models for planning support
Leo G. Kroon, Rob A. Zuidwijk
ERS-2003-032-LIS
<http://hdl.handle.net/1765/332>

How and why communications industry suppliers get "squeezed out" now, and the next phase
L-F Pau
ERS-2003-033-LIS
<http://hdl.handle.net/1765/317>

Financial Markets Analysis by Probabilistic Fuzzy Modelling
Jan van den Berg, Uzay Kaymak, Willem-Max van den Bergh
ERS-2003-036-LIS
<http://hdl.handle.net/1765/323>

WLAN Hot Spot services for the automotive and oil industries :a business analysis or : "Refuel the car with petrol and information , both ways at the gas station "
L-F Pau, M.H.P.Oremus
ERS-2003-039-LIS
<http://hdl.handle.net/1765/318>

A Lotting Method for Electronic Reverse Auctions
U. Kaymak, J.P. Verkade and H.A.B. te Braake
ERS-2003-042-LIS
<http://hdl.handle.net/1765/337>

Supply Chain Optimisation in Animal Husbandry
J.M. Bloemhof, C.M. Smeets, J.A.E.E. van Nunen
ERS-2003-043-LIS
<http://hdl.handle.net/1765/353>

A Framework for Reverse Logistics
Marisa P. de Brito and Rommert Dekker
ERS-2003-045-LIS
<http://hdl.handle.net/1765/354>

An assessment system for rating scientific journals in the field of ergonomics and human factors
Jan Dul and Waldemar Karwowski
ERS-2003-048-LIS
<http://hdl.handle.net/1765/432>

Circulation of Railway Rolling Stock: A Branch-and-Price Approach

Marc Peeters and Leo Kroon

ERS-2003-055-LIS

<http://hdl.handle.net/1765/902>

Emerging Multiple Issue e-Auctions

Jeffrey E. Teich, Hannele Wallenius, Jyrki Wallenius and Otto R. Koppius

ERS-2003-058-LIS

<http://hdl.handle.net/1765/922>

Inventory Management with product returns: the value of information

Marisa P. de Brito and E. A. van der Laan

ERS-2003-060-LIS

<http://hdl.handle.net/1765/925>

Promising Areas for Future Research on Reverse Logistics: an exploratory study

Marisa P. de Brito

ERS-2003-061-LIS

<http://hdl.handle.net/1765/926>

A Polynomial Time Algorithm for a Deterministic Joint Pricing and Inventory Model

Wilco van den Heuvel and Albert P.M. Wagelmans

ERS-2003-065-LIS

<http://hdl.handle.net/1765/929>

A geometric algorithm to solve the $n/g/n/nd$ capacitated lot-sizing problem in $o(t^2)$ time

Wilco van den Heuvel and Albert P.M. Wagelmans

ERS-2003-066-LIS

<http://hdl.handle.net/1765/930>

Arrival Processes for Vessels in a Port Simulation

Eelco van Asperen, Rommert Dekker, Mark Polman, Henk de Swaan Arons & Ludo Waltman

ERS-2003-067-LIS

<http://hdl.handle.net/1765/973>

The distribution-free newsboy problem with resalable returns

Julien Mostard, Rene de Koster and Ruud Teunter

ERS-2003-068-LIS

<http://hdl.handle.net/1765/975>

A note on a multi-period profit maximizing model for retail supply chain management

Wilco van den Heuvel and Albert P.M. Wagelmans

ERS-2003-072-LIS

<http://hdl.handle.net/1765/991>

The Impact of Complexity, Rate of Change and Information Availability on the Production Planning and Control Structure: Evidence from Medium-Sized Dutch Discrete Manufacturing Firms

Marcel van Assen and Steef van de Velde

ERS-2003-083-LIS

<http://hdl.handle.net/1765/1033>

The impact of innovation and organizational factors on APS adoption: Evidence from the Dutch discrete parts industry

Bart van Hezewijk, Marcel van Assen and Steef van de Velde

ERS-2003-084-LIS

<http://hdl.handle.net/1765/1067>

Equilibrium Constrained Optimization Problems

S. Birbil, G. Bouza , J.B.G. Frenk and G. Still

ERS-2003-085-LIS

<http://hdl.handle.net/1765/1068>

Network-Based Business Process Management: Embedding Business Logic In Communications Networks

L-F Pau and P.H.M.Vervest

ERS-2003-086-LIS

<http://hdl.handle.net/1765/1070>