BASICS OF DYNAMIC PROGRAMMING FOR REVENUE MANAGEMENT

Jean Michel Chapuis [*]

Revenue management is the control of inventory and pricing of a perishable product in order to improve the efficiency of its marketing. At the time a customer makes a reservation for a flight, an airline revenue manager must accept the reservation, which is certain at that time, or reject it and wait for a possible demand for seats at a higher rate in the future. Even though this pattern of behaviour has been universal for the last decade, implementation in GDS remains problematic. This article sets out some elements of Dynamic Programming as a new optimisation method. The properties of the opportunity cost of using a unit for a given capacity which is the key to all RM optimisations are here presented in detail.

Le Revenue Management, c'est-à-dire le contrôle de l'inventaire et de la tarification d'un produit périssable, vise à améliorer l'efficience de son marketing. Lors de l'enregistrement des réservations d'un vol, le revenue manager d'une compagnie doit soit accepter une demande de réservation actuelle et certaine ou la rejeter et attendre une probable demande à un tarif plus élevé dans le futur. Alors que sa pratique est omniprésente cette dernière décennie, la mise en œuvre dans les GDS soulèvent de nombreux problèmes. Cet article propose quelques éléments de la Programmation Dynamique comme une nouvelle méthode d'optimisation. Les propriétés du coût d'opportunité de l'utilisation d'une unité pour une capacité donnée, la clé de toute optimisation de RM, sont présentées dans le détail.

I INTRODUCTION

For decades, we have learned to revere the letters of Revenue Management (RM). Its spirit of pricing and of inventory control is widely shared among various industries, major companies and academic researchers. Revenue maximization is usually reached by comparing a current discount

* IRIDIP. Mail: University of French Polynesia, BP 6590, 98702 Faa'a, French Polynesia. Email: jean-michel.chapuis@upf.pf.

Assistant Professor of Management at the University of French Polynesia, PhD (Corporate Finance); Head of the International Hospitality Management Master of the University of La Rochelle, 2002 to 2004; Guest Lecturer for the CHN University of Doha, Qatar and the ESCA of Casablanca, Morocco.

fare request with an expected higher fare request in the future. The well-known EMRS model[1] gives a heuristic solution but assume that the demand is static and will not evolve during the booking period. However, new low cost carrier entrants, fare transparency from Internet fare search engines and changing customer purchasing patterns are the major forces pushing airline revenue managers to adapt their optimization models. In addition, using traditional forecasting and optimization models that assume fare class independence causes continuous erosion of yield. Especially when there is excess capacity, the surplus seats flow down the fare structure, allowing high yield customers to buy down. Later, this behaviour affects forecast, ie expecting lower demand for high fare class in the future, and leads revenue managers to protect less seats for it.

Almost RM situations can usefully be modeled as Dynamic Programming problems (DP) because those models take into account future possible booking decisions in assessing that current decision. Moreover, DP allows to relax the early birds hypothesis (low contribution customers book in advance of the high contribution ones). The optimal controls are then time-dependent as a function of the remaining capacity. Dynamic pricing is the special case where price is the control used by the managers.

Nonetheless, the practising of Bellman equations[2] is relatively new in RM since none of the pioneering papers use DP.[3] McGill and Van Ryzin (1999) conclude that:

> Dynamic formulations of the revenue management problem are required to properly model real world factors like cancellations, overbooking, batch bookings, and interspersed arrivals. Unfortunately, DP formulations, particularly stochastic ones, are well known for their unmanageable growth in size when real-world implementations are attempted. Usually, the only hope for dynamic optimisation in these settings lies in identification and exploitation of structural properties of optimal or near optimal solutions. *Knowledge that an optimal solution must be of a control-limit type or be represented by a*

---

1 See P P Belobaba "Application Of A Probabilistic Decision Model To Airline Seat Inventory Control", *Operations Research* (1989) vol 37, n 2, 183-198; which is an extension of Littlewood's model – K Littlewood "Forecasting and Control of Passengers Bookings" (1972) Proceeds of AGIFORS, Nathanya, Israel.

2 See R Bellman *Applied Dynamic Programming* (Princeton University Press, 1957).

3 The basics of Revenue Management can be found in P E Pfeifer "The Airline Discount Fare Allocation Problem" (1989) *Decision Sciences*, vol 20, n 1, 149-157; Belobaba, op cit or S E Kimes "Yield Management: A Tool for Capacity-Constrained Service Firms" (1989) *Journal of Operational Management*, Vol 8, n 4, 348-64; among others. For an overview, check J I McGill and G R van Ryzin "Revenue Management: Research Overview and Prospects" (1999) *Transportation Science*, vol 33, n 2, 233-257 or G Bitran and R Caldentey "An Overview of Pricing Models for Revenue Management" (2003) *Manufacturing and Service Operations Management*, vol 5, 202-229, and more recently W-C Chiang, J C H Chen and X Xu "An Overview of Research on Revenue Management: Current Issues and Future Research" (2007) *International Journal of Revenue Management*, vol 1, n 1, 97–128. The experience of American Airlines while implementing Revenue Management in a worldwide scale is reported in B C SMITH, J F LEIMKUHLER AND R M DARROW "Yield Management at American Airlines" (1992) *interfaces*, vol 22, n 1, 8-31.

> *monotonic threshold curve can be invaluable in development of implementable systems*. The existing literature has already identified such structures in special cases of the revenue management problem; however, there are difficult areas still requiring work (emphasis added).

The goal of this research is to review some Dynamic Programming models dedicated to Revenue Management in order to provide a solid basis for future work. The structure is the following: Part II, formulates both the RM problematic and its DP methodology. Then, the problem of setting booking limits and prices using DP approach is addressed in the Part III. The last Part concludes and opens applications of DP in RM.

II DYNAMIC PROGRAMMING OF REVENUE MANAGEMENT OPTIMISATION

After reviewing the RM problem, a brief review of the literature of DP models applied to RM is presented in §A and a mathematical formulation is in §B. As an example, consider the following case. You are managing a hotel of two rooms for the next weekend (the Friday and Saturday nights), meaning you have to sell 4 units of a same resource, marketed through 3 products. You charge $100 for a room per night, except for the guests who stay for 2 consecutive nights. The latter costs $160. This is Monday and you check on your web site the actual activity in order to confirm reservation. There are 2 requests for the Friday night (at $100 each), 1 request for the Saturday night (still at $100). At the same time, your travel agent calls you to book 2 couples for 2 nights (so the price would be $160 each). What should you do? Revenue Management can help to choose which requests to accept.

In order to maximise operating revenue, RM is the process by which a manager controls the availability of a product or service, marketed with a differential and dynamic pricing. Controls can be effective by varying prices, setting booking limits and managing fences.[4] This approach is in line with Talluri and Van Ryzin (2004),[5] who supported both price and quantity-based models of RM. Most authors often define RM as the application of control and pricing strategies to sell the right capacity to the right customer, in the right place, at the right time and at the right price[6]. RM has proven its potential impact on profitability in the past[7].

---

4 In the previous example, one can easily see that the first come first serve rule is not optimal for the hotel. The fence of "stay at least two nights", that can justify the discount rate and help to increase revenue, is not totally efficient. Only a booking limit of one "Two nights" package or a bid price for each night are optimal.

5 See K T Talluri and G J van Ryzin "The Theory and Practice of Revenue Management" 712.

6 Kimes, op cit.

7 Smith, Leimkulher and Darrow, op cit.

A Literature review of DP Models to RM Optimisation.

Dynamic programming addresses how to make optimal decisions over time under uncertain conditions and to control a system[8]. Most RM situations can be analysed assuming a discrete-state and a discrete time over a finite-horizon modelling[9]. Table 1 summarises the literature according to four criteria: (i) A paper could consider a single product (at various prices) or multiple products (depending on purchase restrictions or independent demands for example); (ii) A paper could consider a static policy (assuming a strict order of booking arrivals) or allow for a dynamic policy (not assuming the early birds hypothesis); (iii) A paper could consider various forms of demand process; (iv) A paper could consider either a single resource for 1 to *n* products or multiple resources (such as an airline network of hubs and spokes).

Table 1: classification of literature

	Product
Single	Curry (1990); Wollmer (1992); Brumelle and McGill (1993); Lee and Hersh (1993); Gallego and Van Ryzin (1994); Bitran and Mondshein (1995); Robinson (1995); Lautenbacher and Stidam (1999); Zhao and Zheng (2000); You (2001)
Multiple	Gallego and Van Ryzin (1997); Talluri and Van Ryzin (1998); Feng and Xiao (2001); Kleywegt (2001); Bertsimas and Popescu (2003); El-Haber and El-Taha (2004); Bertsimas and De Boer (2005)
	Policy
Static[10]	Curry (1990); Wollmer (1992); Brumelle and McGill (1993)

8 The methodological reference is D Bertsekas *Dynamic Programming and Optimal Control* (1995) Athena Scientific, vol 1, Belmont MA.

9 T Lee and M Hersh "A Model for Dynamic Seat Inventory Control with Multiple Seat Bookings" (1993) *Transportation Science*, vol 27, pp. 233-247 is the most popular reference with their model for dynamic seat inventory control. G Gallego and G van Ryzin "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons" (1994) *Management Science*, vol 40, n 8, 999-1020 are commonly cited as the reference in dynamic pricing model using dynamic programming.

10 D Bertismas and S de Boer "Dynamic pricing and inventory control for multiple products" (2005) *Journal of Revenue and Pricing Management* vol 3, n 4, 303-319; D Bertismas and I Popescu "Revenue Management in a Dynamic Network Environment" (2003) *Transportation Science*, vol 37, n 3, 257-277; G Bitran and S Mondschein "An Application of Yield Management to the Hotel Industry considering Multiple Day Stays" (1995) *Operations Research*, vol 43, n 3, 427-443; S Brumelle and J I McGill "Airline Seat Allocation with Multiple Nested Fare Classes" (1993) *Operations Research* 137; R Curry "Optimal Seat Allocation with Fare Classes Nested by Origins and Destinations" (1990) *Transportation Science* vol 24, 194-204; Y Feng and G Gallego, 2000, «Perishable asset revenue management with Markovian time dependent demand intensities», *Management Science*, vol 46, n 7, 941-957; Y Feng and B Xiao, 2001, «A dynamic airline seat inventory control model and its optimal policy» *Operations Research*, vol 49, n 6, 938-952; G Gallego and G van Ryzin "A Multiproduct Dynamic Pricing Problem and its Applications to

Dynamic	Lee and Hersh (1993); Gallego and Van Ryzin (1994); Robinson (1995); Gallego and Van Ryzin (1997); Talluri and Van Ryzin (1998); Liang (1999); You (1999); Lautenbacher and Stidam (1999); Feng and Gallego (2000); Feng and Xiao (2001); Kleywegt (2001); You (2001); Bertsimas and Popescu (2003); El-Haber and El-Taha (2004)
Both	Bertsimas and De Boer (2005)
	Demand
Deterministic	Curry (1990); Wollmer (1992); Brumelle and McGill (1993); Robinson (1995); Kleywegt (2001)
Stochastic	Lee and Hersh (1993); Gallego and Van Ryzin (1994); Gallego and Van Ryzin (1997); Talluri and Van Ryzin (1998); Feng and Gallego (2000); Feng and Xiao (2001); You (2001); Bertsimas and Popescu (2003); El-Haber and El-Taha (2004)
Non homogeneous	Zhao and Zheng (2000); Bertsimas and De Boer (2005)
	Resources
Single	Wollmer (1992); Brumelle and McGill (1993); Lee and Hersh (1993); Gallego and Van Ryzin (1994); Robinson (1995); Lautenbacher and Stidam (1999); Feng and Gallego (2000); Zhao and Zheng (2000); Kleywegt (2001); You (2001); Bertsimas and De Boer (2005)
Network	Curry (1990); Bitran and Mondshein (1995); Gallego and Van Ryzin (1997); Talluri and Van Ryzin (1998); You (1999); Feng and Xiao (2001); Bertsimas and Popescu (2003); El-Haber and El-Taha (2004)

The way the behaviour of customers is incorporated in the optimisation process is the next challenge. The following part almost borrows from Talluri and Van Ryzin book.[11]

Network Yield Management" (1997) *Operations Research*, vol 45, n 1, 24-42; S E Kimes "Yield Management: A Tool for Capacity-Constrained Service Firms" (1989) *Journal of Operational Management*, Vol 8, n 4, 348-364; A J Kleywegt "An Optimal Control Problem of Dynamic Pricing" (2001) *Working paper*, 25; C Lautenbacher and S Stidham "The Underlying Markov Decision Process in the Single-Leg Airline Yield Management problem" (1999) *Transportation Science*, vol 33, n 2, 136-146; L W Robinson "Optimal and Approximate Control Policies for Airline Booking with Sequential Fare Classes" (1995) *Operations Research*, vol 43, n 1, 252-263; K T Talluri and G J van Ryzin «An aAnalysis of Bid-Price Controls for Network RM» (1990) *Management Science*, vol 44, n 11, 1577-1594; R Wollmer "An Airline Seat Management Model for a Single Leg Route when Lower Fare Classes Book First" (1992) *Operations Research*, vol 40, 26-37; P S You "Airline Seat Management with rejection-for-possible-upgrade decision" (2001) *Transportation science part B*, vol 35, 507-524, W Zhao and Y S Zheng "Optimal Dynamic Pricing for Perishable Assets with Nonhomogeneous Demand" (2000) *Management Science*, vol 46, n 3, 375-388.

11 2004, Ibidem, 651

B A Simple Formulation of RM with a Dynamic Program

This paper choose a traditional RM problem by considering $n > 2$ fare classes. Demand for a single product is then considered discrete as well as capacity (assuming a single resource to simplify). Dynamic models do not assume the traditional hypothesis of early birds. Demand can arrive in a non-strict increasing order of revenue values. Demand is assumed independent between classes and over time and also independent of the capacity controls. The n classes are indexed by j such that $p_1 > p_2 > \ldots > p_j > \ldots > p_n$. Multiple booking is not considered.[12]

Over time, this *system* evolves as a function of both *control decisions* and *random disturbances* according to a *system equation*. The system generates rewards that are a function of both the state and the control decisions. The objective is to find a *control policy* that maximises the total expected revenues from the selling period. There are T time-periods. Time, indexed by t, runs reverse so that $t = 1$ is the last period and $t = \mathrm{T}$ is the first period. By the way, t indicates the time remaining before the product is consumed or loses its value. Applying basics of DP to RM gives the following:

- $\mathbf{x}(t)$ is, as the state of a system, the remaining capacity constrained between 0 and the full capacity C (both cancellation and overbooking are not considered).
- $\mathbf{w}(t)$ is the random disturbance representing the consumer's demand to the firm for product j. The probability of an arrival of class j in the period t is λ_t^j . The horizon is divided into decision periods that are small enough so that no more than one customer arrives during each period, thus $\sum_{j=1}^{n} \lambda_t^j \leq 1$.
- $\mathbf{u}(t)$ is the control decision, assumed discrete and constrained to a finite set that may depend on time t and the current state $\mathbf{x}(t)$. $\mathbf{u}(t)$ is then the quantity u of demand to accept. In the simplest statement, u is constrained to be either 0 or 1. The amount accepted must be no more than the capacity remaining, so $u \leq x$.

A few points must be noted. First, as the manager waits for the demand to realise to make his or her decision - accept or reject an incoming booking request - the random disturbance is observable. In other words, s/he can build the control action with perfect knowledge of the disturbance.[13]

Second, customers are allowed to arrive in any order (interspersed arrivals), in contrast to static models that assume sequential booking classes or low-before-high fares. Accepting this early-birds hypothesis leads to think in terms of the *number of bookings* asked per period (ie the booking limit)

---

12 Usually, researchers consider that the time can be divided in small enough intervals such that only one booking arrives. You, Op cit develops a model allowing for multiple bookings.

13 This assumption allows to greatly simplify both the mathematical formulation and the optimal control policies because the optimisation problem could be spread in multiple subperiods optimisation problems that are independents between them. See Talluri and Van Ryzin (2004, *ibid*, 654) to understand the full consequence of this assumption.

and only requires forecasting the total demand for each class, D_j. However, the optimal control would be "static", meaning constant over time (how much to protect or when to reject booking) as the manager never came back on his or her decision. Another approach would have been to reason for *only one booking request* at a time. The incoming requests are not ordered by price but one needs to model the demand process and then the optimal control changes every period according to the remaining capacity – so "dynamic" means real time adjustment.

- $\mathbf{g}_t(\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t))$ is a real-valued reward function, specifying the revenue in period t as a function of the parameters. The total revenue is additive and terminal one is assumed to be zero whatever happens; namely that the remaining capacity after the last stage is no value.

The objective is to maximise the total expected revenue for T periods $E\left[\sum_{t=1}^{T} g_t\left(x(t),u(t),w(t)\right)\right]$ by choosing T control decisions u_T, …, u_2, u_1. This is a Markovian control because the control depends only on the current state and the current time, and no other information is needed, such as the history of the process up to time. The collection $\{u_T, \ldots, u_2, u_1\}$ is called a policy μ.

For a given initial state $\mathbf{x}(T) = x$, the expected (total) revenue of a policy is:

$$V_T^{\mu}(x)=E\left[\sum_{t=1}^{T} g_t\left(x(t),u_t(x(t)),w(t)\right)\right]$$

In other words, this is the total expected revenue that can be generated when there are T decision periods (remaining) and x products to sell. E[] means the expectation is estimated over j, the fare classes or products. The optimal policy, denoted u^*, is one that maximizes $V_T^{u^*}(x)$ and is simply written $V_T(x)$. The principle of optimality[14] lies at the heart of DP. It says that if a policy is optimal for the original problem, then it must be optimal for any subproblem of this original problem as well. Let say that if $\{u_t^*, u_{t-1}^*, \ldots, u_1^*\}$ is not optimal for the t-subproblem and another policy, $\hat{u}$, yields a greater expected reward, then the optimality of u^* is contradicted. Because the policy $\{u_T^*, \ldots, u_{t+1}^*, \hat{u}_t, \ldots, \hat{u}_1\}$ would produce a strictly greater expected revenue than does the policy u^*.

Applying this principle leads to use a recursive procedure for finding the optimal policy. The following value function $V_t(x)$ is the unique solution to the recursion shown in equation 1, for all t and all x.

Once the value of the demand is observed, the value of u is chosen to maximize the current period t revenue plus the revenue to go, or $p_j.u + V_{t-1}(x - u)$ subject to the constraint $u \in \{0,1\}$. This sentence leads to think in terms of $V(x) = E[\max\{\}]$ instead of $V(x) = \max\{E[]\}$ as in traditional DP, due to the previous assumption about the random disturbance. The value function entering period t, $V_t(x)$, is then the expected value of this optimization (namely the maximum) with respect of the demand and is given in the form of the Bellman equation:

Equation 1: $$V_t(x) = E\left[\max_{u\in\{0,1\}}\left\{p_j u + V_{t-1}\left(x-u\right)\right\}\right]$$

---

14 Due to Bellman (1957), op cit.

With the boundary $V_0(x) = 0$, whatever is x, and $V_t(0) = 0$, whatever is t. This means that the unsold inventory left the last day is useless and a sunk cost.

The motivation of Equation 1 uses the Bellman's principle of optimality. Since $V_{t-1}(x - u)$ is the optimal expected value of the future revenue given the state $(x - u)$ in the next period, t -1, the optimal value of the t-subproblem should be the result of maximizing the sum of the current expected reward $E[p_j u]$ and the expected reward for the t - 1 subproblem, $E[V_{t-1}(x - u)]$. This is the optimal solution for the period t and the process can be restarted until reaching T.

III APPLICATION OF DP TO FIND OUT BOOKING LIMITS AND PRICING ISSUES

The RM problem is basically formulated as a marginal one, such that the opportunity cost or expected marginal value (hereafter EMV) becomes the cornerstone of any model. The properties that any control must hold to be optimal are shown in §A. Nevertheless, there exists various controls, presented in §B, as managers in practice figure out few control types such as a threshold. There is also a close correspondence between DP model and Belobaba's EMSR model.

A The expected marginal value as the key solution

The values u^* that maximise the right-hand side of equation 1 for each t and x form an optimal control policy for this model. The solution can be found by evaluating the EMV at period t of the x^{th} unit of capacity.

Equation 2.: $\Delta V_t(x) = V_t(x) - V_t(x-1)$

The most important concerns are how this marginal value behaves with changes in the capacity left x and the t remaining periods. The expected marginal value of $V_t(x)$ satisfies[15] $\forall x, t$:

(i) $\Delta V_t(x+1) \leq \Delta V_t(x)$

EMV is decreasing in x, meaning that the opportunity cost of using one more unit of the capacity increases when there are less products available. In other words, the fewer products on hand the more the chance to loose a sale.

(ii) $\Delta V_t(x) \geq \Delta V_{t-1}(x)$

EMV is increasing in time period left or is decreasing as time elapses, meaning that the opportunity cost of using one more unit of the capacity is higher when the time to go is long than

15 There are different proofs of these statements. However, the Talluri and Van Ryzin (2004, ibid, 38) proof is one of the most clever. Hence, the seat allocation problem and the dynamic pricing problem satisfy well-known sufficient conditions for an optimal policy to be monotonic. These problems translate to the existence of time-dependent controls. In other words, the optimality value function is concave and non-increasing, from which it follows that an optimal admission policy is monotonic in the state.

when it's short. In other words, the longer the remaining time the more the chance to use the capacity efficiently (the probability of a customer requesting a full fare is high).

These two properties are intuitive and greatly simplify the control decisions because the equation 1 can be rewritten with equation 2.

Equation 3: $$V_t(x) = V_{t-1}(x) + E\left[\max_{u \in \{0,1\}} \left\{p_j u - \Delta V_{t-1}(x-u); 0\right\}\right]$$

The meaning of this equation is the following. The value of the revenue to go today (in t) is equal to the revenue to go tomorrow (that is already optimised using a recursive method) plus the expected maximum revenue of the decision to make (that is the fare net of the opportunity cost if positive). The problem is formulated as a backward-recursion dynamic program. In other words, an optimal booking policy is reached by the assessment of accepting a booking request relative to the decrease in expected total revenue associated with removing one product from the available inventory.

B Types of Controls

Three types of controls can be derived, studying the sign of value of the max function. Needless to say that relaxing the hypothesis of early birds implies that the protection levels, the booking limits and the bid prices are time-dependent because the random disturbance associated with the value function depends on the probability of an arrival of class j in the period t. In order to an efficient optimization, the manager must know the process of demand for each fare class over time (*ie* the booking curve).

1 Protect levels and nested booking limits

First consider controlling the revenue through u, the accept (u = 1) or reject (u = 0) decision alternatives to the current request. Since $\Delta V_t(x)$ is decreasing in x, it follows that $p_j u - \Delta V_t(x - u)$ is increasing in x. For a given t, it is optimal to keep accepting incoming requests until the previous term becomes negative or the upper bound $\min\{D_j; x\}$ is reached, whichever comes first.

The protection level y_j^* is the number of units to save for customers who request, or will request in the future, class j or higher (*ie* j, j - 1, …, 1). It is the maximum quantity x such that the fare level is too low to compensate the opportunity cost, justifying the reject decision (see equation 4).

Equation 4: $y_j^* = \max\ \{x: p_{j+1} < \Delta V_t(x)\}$, j = 1,…, n-1.

The optimal control at the period t + 1 is given by equation 5:

Equation 5: $u^*(t+1; x; D_{j+1}) = \min\{(x - y_j^*)^+; D_{j+1}\}$.

The booking limit b^*_{j+1} is the quantity is excess of the protection level ($x - y_j^*$) if positive or zero otherwise. This is the maximum space available for class j + 1 booking request. Those booking limits are nested, meaning that a class j request can be withdrawn in any of the j class and above (to n). By the way, this set of protected classes (j, j - 1, …, 1) are proposed or "open" to customers being denied a request fare class below j. The policy then is "simply accept requests first come, first serve until the capacity threshold b^*_j is reached or the time ends, whichever come first. The figure 1 represents the accept/reject alternatives of a decision. For an illustration, consider in the figure 2 where $n = 3$ and $p_1 > p_2 > p_3$.

Figure 1: critical booking capacity

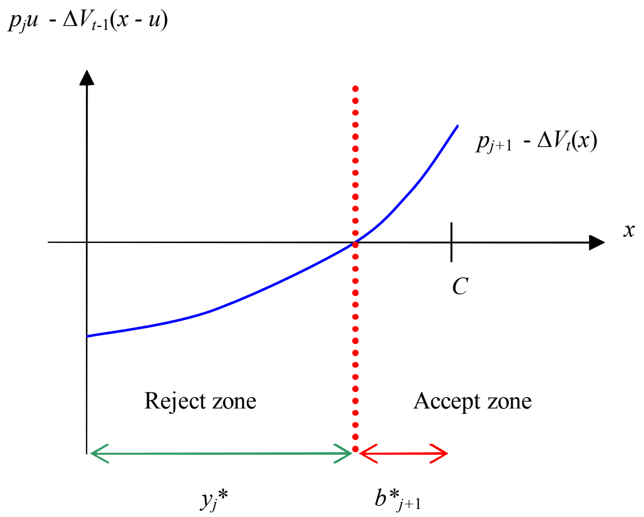

Legend: for a given t, the decision to accept or reject a request with fare j+1 depends on x the level of capacity left. The critical booking capacity separates the accept zone and the reject zone.

2 The control by time thresholds

Second, consider controlling the revenue through t, the time left, to accept (u = 1) or reject (u = 0) an incoming request. Based on previous properties of equation 2, one can show that time thresholds characterize the optimal policy. During the booking horizon, they are points in time before which requests are rejected and after which requests are accepted. Since $\Delta V_t(x)$ is increasing in t, it follows that $p_j u - \Delta V_{t-1}(x - u)$ is decreasing in t. Thus, it is optimal to keep refusing incoming requests until the previous term becomes positive (see figure 3).

This proposition is also easy to understand with figure 1. Since $\Delta V_t(x)$ is increasing in t, one should accept more booking for a given j when remaining time lessens. In other words, as time elapses, the EMV decreases and the line lifts up. Then the optimal protection level y^*_j will shift to the left (decrease). The nested protection structure is kind of $\ldots < y^{*t-1}_j < y^{*t}_j < y^{*t+1}_j < \ldots$

Figure 2: Case with 3 fare classes

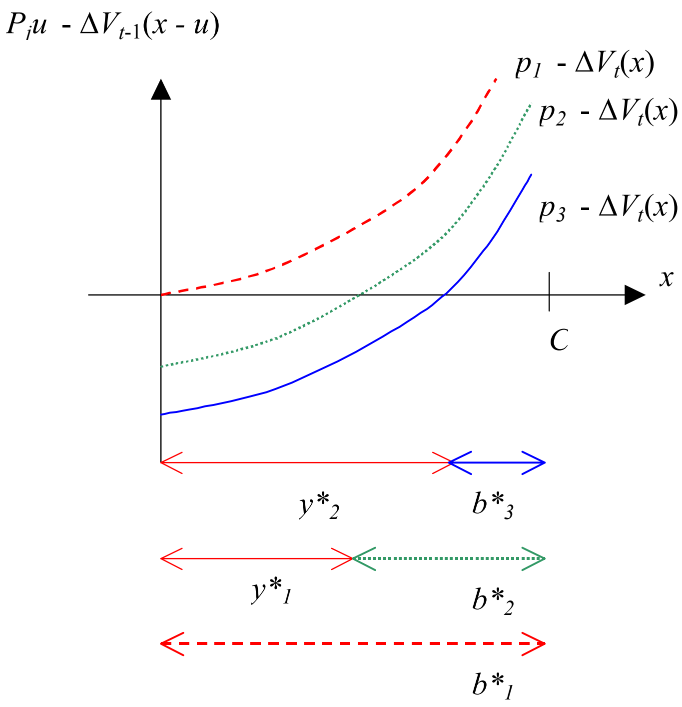

3 *Bid prices*

Finally, the optimal control can also be implemented through a table of bid prices, defined as $\pi_{j+1}(x) = \Delta V_t(x)$. A bid price is the minimum amount of money to accept in exchange of a unit of capacity. Under mild assumptions (Gallego and Van Ryzin, 1994; Talluri and Van Ryzin, 1998), the optimal price path $\pi_j(x)$ is decreasing in x and increasing in t. Dynamic pricing is becoming a real challenge for airlines since low cost carriers are introducing less restricted fares. For example, Ryanair or Easyjet are not able to be sure that the more price sensitive passengers (low fare) book first. Then they apply a dynamic pricing strategy for a single product without price discrimination between their customers at a given time.

Figure 3: critical time threshold

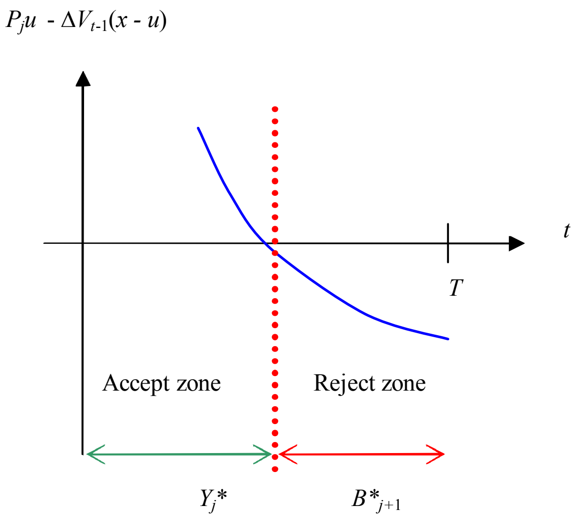

Legend: For a given booking capacity x, a request for a seat of fare class j in decision period t is accepted if time remaining is short (less than Y_j*, the critical time threshold) and rejected otherwise. The caps denote that the control is expressed on a time axis instead of a quantity axis.

IV DISCUSSION AND CONCLUSION

The use of Dynamic Programming in Revenue Management helps to decide whether to accept or reject an incoming booking reservation with more realism than older methods. There are two main points. One is the proposition that DP in RM allows to relax the low-before-high fare order of arrival bookings. In practice, the DP provides the optimal policy for the RM problem, by evaluating the whole tree of possibilities and making at each point in time the decision that would imply higher future expected revenues, processing backward recursion. The dark side is the increase in the computation difficulties according to the dimension of the problem. This means that for a single product with 100 units to sell over a 200 time periods, the number of iterations is 100 x 200 = 20 000. But for three products using the same resource, this number becomes 100 x 200^3 = 800 millions. United Airlines expects a 1 to 2 % increase in revenue ($ 158 billions of 2005 turnover) due to the implementation of DP[16].

16 K Saranathan and W Zhao "Revenue Management in the New Management of the New Fare Environment " (2005) Proceeds of AGIFORS Reservation and Yield Management Study Group, Cape Town.

The second interest of DP approach, and the main avenue for future research, is that it allows RM to incorporate consumer choice within the optimisation process. El-Haber and El-Taha[17] formulate a dynamic programming model to solve the seat inventory control problem for a two-leg airline with realistic elements of consumer behaviour. Ahead of the Origin and Destination formulation, they consider cancellation, no-shows and overbooking. Following Talluri and Van Ryzin's work,[18] Van Ryzin and Vulcano[19] consider a revenue management, network capacity control problem in a setting where heterogeneous customers choose among the various products offered by a firm (for example, different fight times, fare classes and/or routings). Customers may therefore substitute if their preferred products are not offered, even buy up. Their choice model is very general, simply specifying the probability of purchase for each fare product as a function of the set of fare products offered. Overall, the value of this paper is to facilitate the understanding of more complex, and probably more realistic, models of revenue management.

---

17 S El-Haber and M El-Taha "Dynamic two-leg airline seat inventory control with overbooking, cancellations and no-shows" (2004) *Journal of Revenue and Pricing Management*, vol 3, n 2, 143-170.

18 K T Talluri and G J van Ryzin "Revenue Management under a General Discrete Choice Model of Consumer Behavior" (2003) *Management Science*, vol 50, n 1, 15–33.

19 G J van Ryzin and G Vulcano "Computing Virtual Nesting Controls for Network Revenue Management Under Customer Choice Behavior" (2006) *Columbia University working paper*.