

**Open**

Risk minimising strategies for revenue management problems with target values

Matthias Koenig^{1*} and Joern Meissner²¹Bielefeld University of Applied Sciences, Minden, Germany; and ²Kuehne Logistics University, Hamburg, Germany

Consider a risk-averse decision maker in the setting of a single-leg dynamic revenue management problem with revenue controlled by limiting capacity for a fixed set of prices. Instead of focussing on maximising the expected revenue, the decision maker has the main objective of minimising the risk of failing to achieve a given target revenue. Interpreting the revenue management problem in the framework of finite Markov decision processes, we augment the state space of the risk-neutral problem definition and change the objective function to the probability of failing a certain specified target revenue. This enables us to obtain a dynamic programming solution that generates the policy minimising the risk of not attaining this target revenue. We compare this solution with recently proposed risk-sensitive policies in a numerical study and discuss advantages and limitations.

Journal of the Operational Research Society (2016) 67(3), 402–411. doi:10.1057/jors.2015.63

Published online 2 September 2015

Keywords: capacity control; revenue management; multi-period; risk; target level criterion

The online version of this article is available Open Access

1. Introduction

Revenue management systems have become a standard tool in various industries beyond the original airline industry. These industries range from cruise lines, rental cars, media advertising, medical services to event management (see eg Talluri and van Ryzin, 2005).

We consider a typical revenue management model: a firm operating in a monopolistic setting offering multiple products. These products consume a fixed resource of a limited capacity. The firm sells the products over a finite time horizon. At the end of this time, the salvage value of the resource is assumed to be 0.

The firm can influence its revenue stream by allocating capacity to different classes of demand. Its objective is to find a policy which optimises an objective function. Normally, this objective function is risk-neutral, and the policy is chosen to maximise expected revenue. Such a risk-neutral objective can be motivated by the law of large numbers if the revenue process repeats itself very often, for example, a daily operating airline flight connection.

However, a risk-neutral policy might not be requested under all scenarios and a risk-averse policy might be advantageous for the decision maker. Lancaster (2003) remarks that a risk-neutral model is often not sufficient, even in the airline industry, as a stable revenue might be preferable because of financial constraints.

In practice, decision makers present some level of risk aversion in revenue management, as mentioned by Bitran and Caldentey (2003). Weatherford (2004) reports the same experience. He observed that airline analysts feel uncomfortable with recommendations of their (risk-neutral) revenue management systems, in particular while waiting for the high-fare passengers a few days before flight departure.

In recent papers by Barz and Waldmann (2007), Huang and Chang (2011) and Koenig and Meissner (2015), risk-neutral and risk-sensitive policies are analysed. The results show that an appropriate risk-averse policy can be selected if the decision maker knows the parameters representing his level of risk aversion. Such parameters have to be determined, which is something that is not straightforward in either of the published approaches, whether the underlying concept is an exponential utility function or a discount factor relaxing an optimality condition. Usually, the parameters have to be estimated by running numerical experiments and evaluating risk measures, such as mean-variance or conditional-value-at-risk, on the results.

Thus, we propose using the target-percentile risk measure, discussed by Boda and Filar (2006), as the object function. The target-percentile risk measure computes the probability of the return failing to achieve a previous given fixed target. There are several advantages of using this measure. First, one important structural property is its time consistency. It says that optimality of decisions should only consider the future. Time consistency is a desirable property for multi-period risk measures as it allows its use in dynamic programming, as shown for example by the work of Shapiro (2009). Second, it does not assume a special kind of revenue distribution, as it measures the

*Correspondence: Matthias Koenig, Bielefeld University of Applied Sciences, Artilleriestr. 9, 32427 Minden, Germany.
E-mail: matthias.koenig@fh-bielefeld.de

percentile of the given target level. Third, numerical computation schemes are available as described by Wu and Lin (1999). Fourth, Boda and Filar (2006) show that multi-stage versions for the well-established risk measure value-at-risk can be developed using the target-percentile measure.

Fifth and important, it is easily interpreted by practitioners and does not require a risk sensitivity parameter which might be difficult to assess. Practitioners know the cash constraints of their businesses, which enable them financial liquidity and operational freedom. For their businesses, thus, they know the desired target level which they can use as input parameter in our model.

The structure of the paper is as follows. We review the relevant literature in Section 2. In Section 3, we describe our model as a Markov decision process and its extension to apply the target-percentile risk measure. This section also contains some implementation details. Section 4 shows numerical results of our approach and provides a comparison with results of other approaches. Finally, we conclude the paper in Section 5.

2. Related work

Most revenue management models use a risk-neutral objective function. We refer to the work of Talluri and van Ryzin (2005) for an overview of these kinds of models. In general, revenue management models are categorised often as capacity control model or dynamic pricing. However, Maglaras and Meissner (2006) discuss the similarities between both categories and give a common formulation in a risk-neutral setting.

The risk-neutral model of dynamic capacity control, which we consider here, was introduced by Lee and Hersh (1993). The corresponding Markov decision process is described by Lautenbacher and Stidham (1999).

The approaches for incorporating risk in revenue management models are analogous to the general decision making under risk: expected utility theory, mean-variance considerations, probabilistic constraints.

Expected utility theory as an element for reflecting risk in revenue management is recommended by Weatherford (2004). He states that the assumption of risk neutrality is not given for many practical scenarios and proposes expected utility theory as a risk-averse solution. Instead the well-adopted (risk-neutral) expected marginal seat revenue model, standard algorithms introduced by Beloba (1989), the expected marginal seat utility heuristic can reflect risk sensitivity for decision making. Weatherford and Beloba (2002) also show how forecasting errors affect the revenue.

Recent works of Barz and Waldmann (2007), Feng and Xiao (2008) and Xiong *et al* (2011) are employing expected utility theory, too. Both papers support the application of an exponential utility function to account for risk aversion. Barz and Waldmann (2007) use the Markov decision process formulation of static and dynamic capacity control models, whereas Feng and Xiao (2008) provide closed form solutions from a more

general point of view, and Xiong *et al* (2011) consider overbooking in their model.

As the first revenue management model with risk considerations, the model of Feng and Xiao (1999) uses variance as its risk measure; in particular, the variance of sales because of price changes. In order to integrate risk into their objective function, they combine expected revenue with a weighted penalty function for the sales variance. The risk sensitivity of the decision maker can be adjusted by the weighting.

Recently, Huang and Chang (2011) presented a risk-sensitive modification of the optimality condition for the dynamic capacity control model and investigated their method by measuring mean *versus* standard deviation in simulation runs. They offer a ranking of their risk-sensitive policies using a Sharpe ratio of revenue and standard deviation.

Illustrating the vulnerability of risk-neutral revenue management because of demand forecast inaccuracy, Lancaster (2003) recommends a revenue-per-available-seat-mile-at-risk metric, which integrates risk measurement with the value at risk (V@R) metric. This metric is the expected maximum of underperformance over a time horizon at a chosen confidence level.

That cost of price changes should be considered from a risk perspective is demonstrated by Koenig and Meissner (2010) who compare the suitability of two different pricing strategies by the risk measures standard deviation and conditional value at risk (CV@R). In a further paper, Koenig and Meissner (2015) evaluate a range of risk-sensitive policies for the dynamic capacity control model. Gönsch and Hassler (2014) propose an heuristic for computing an CV@R-optimal policy in a recent paper. Their approach solves a knapsack problem for each state in their value function.

Risk sensitivity is incorporated by Levin *et al* (2008) into a dynamic pricing model of perishable products. Their objective function consists of maximum expected revenue constrained by a desired minimum level of revenue with minimum acceptable probability. This constraint is similar to a V@R formulation. The authors formulate a hybrid objective function which combines the risk-neutral objective of expected revenue and a penalty term representing risk aversion. Principally, they approach this risk-adjusted maximisation problem by using a further state which a risk-neutral dynamic pricing model does not require. This state keeps track of already gained revenue.

In a capacity control setting, we base our risk incorporation on a state space expansion, too, but our underlying model is derived from a Markov decision process formulation.

3. Description of model

In the following, we describe the dynamic capacity control problem as a Markov decision process in a similar way as previously done by Lautenbacher and Stidham (1999) and Barz and Waldmann (2007). This model is then expanded in the state space in order to become a model which allows the application of a risk-minimising policy. To achieve this, we follow the

approach of Wu and Lin (1999). Our objective function is the target-percentile dynamic risk measure. Finally, we point out some aspects for implementation of this approach.

3.1. Markov decision process for dynamic capacity control model

We consider the capacity control model stated by Lee and Hersh (1993), which is often referred to as dynamic capacity control. Although originally developed for airline revenue management, it can be transferred to other industries. We describe the model in terms of its original airline revenue management context in order to be more intuitive.

We assume that the booking requests follow a Poisson arrival process. Thus, the booking period for a single-leg flight is separated into N decision periods in such a way that the probability of more than one request can be ignored. The decision periods are denoted by $n \in \{0, \dots, N\}$. The departure is at $n = 0$. If it supports understanding, we will use n as subscript else omit it. Further, there are k booking classes with fares $F_i, F_1 > F_2 > \dots > F_k > 0$ and $F = \{F_1, \dots, F_k\}$. The probability of a request for fare class i in decision period n is given by $p_{n,i}$. Further, we set the probabilities for $n = 0$ as zero for all fare classes: $p_{0,i} = 0$; this step just supports our model setting as the last decision will be made at time $n = 1$. The probability of no request in period n is $p_{n,0} = 1 - \sum_{i=1}^k p_{n,i}$. The initial capacity of seats is given by C . The remaining seats are given by $c \leq C$ in a time period.

We have a finite-state, discrete-time, Markov decision process $\Gamma = (S, A, R, P)$ with state space S and action space A . Further, R denotes the reward set and P , the set of transition probabilities. Time runs in discrete steps and represents the remaining time before flight departure.

The state space S contains all possible configurations of remaining capacity c and request for a fare class i . Thus $S = \{0, 1, \dots, C\} \times \{0, 1, \dots, k\}$ and a state $(c, i) \in S$ says that we have c seats left and a request for fare class i . We set the fare class 0 with fare $F_0 = 0$, as is often common practice.

Our action space $A(c, i)$ corresponds to the ‘reject’ and ‘accept’ decisions for a given state. We have $A(c, i) = \{0, 1\} \forall (c, i) \in S, i > 0$ and $A(c, 0) = A(0, i) = \{0\}$ to only allow the accepting and rejecting of seats at the valid fare prices and not for the artificial class $i = 0$. Overbooking is not allowed.

Let R be the set of rewards (fares) when accepting one booking. Rewards are denoted by $r_n(s, a) \in R$ with $s \in S, a \in A$ and $r_n((c, i), a) = aF_i$ for $n, c > 0$ and 0 otherwise. The transition probabilities $p \in P$ are defined for states $(c, i), (c - a, j) \in S$ with $a \in A$ by $p_n((c - a, j)|(c, i), a) = p_{n,j}$ for $n = N, N - 1, \dots, 0$, and 0 otherwise.

A decision maker decides on a sequence of rules $a_n = d_n(c_n, i_n)$, which determine a policy $\pi = \{d_N, d_{N-1}, \dots, d_1\}$. Thus, a policy determines if a booking request is accepted or rejected in state (c_n, i_n) .

Now let $\rho_N^\pi(c, i) = \sum_{n=0}^N r_n$ denote the random variable of the gained revenue for a particular policy π beginning with

capacity c and request i at N remaining time steps. The expected revenue is given by

$$v_N^\pi(c, i) = \mathcal{E}_\pi[\rho_N^\pi(c, i)] = \mathcal{E}_\pi\left[\sum_{n=1}^N r_n((c_n, i_n), d_n(c_n, i_n))\right].$$

The maximal expected revenue and its associated risk-neutral policy can be computed by the Bellman equation for this problem. However, we are interested in a policy which minimises the time-consistent dynamic risk measure of not achieving a target revenue x in the accumulated return.

3.2. Markov decision process for minimising risk of failing target

We are interested in minimising the risk of not attaining a specified target revenue x_N for the dynamic capacity control model. Thus, we want to find a policy π , which minimises the objective function representing the probability of not achieving a previous specified target level x_N . In order to derive this objective function, we follow the approaches mentioned by White (1988), Wu and Lin (1999) and Boda and Filar (2006) and expand the Markov decision process Γ by a larger state space. The extended Markov decision process $\tilde{\Gamma}$ is similar to Γ . It consists of $\tilde{\Gamma} = (\tilde{S}, \tilde{A}, \tilde{R}, \tilde{P}) = (\tilde{S}, A, R, P)$, as described below.

The state space S is replaced by the new state space $\tilde{S} = S \times \mathbb{R}$ with elements $((c, i), x)$ with $\mathbb{R} = (-\infty, \infty)$. The new state space \tilde{S} consists of states of the configurations of remaining capacity c and a request for fare class i , and additionally, a revenue target x . All state variables are updated over time, for example, the revenue target x decreases by the realised fare price in accordance with decrementing c by selling a seat.

The action space \tilde{A} is generated from action state A by $\tilde{A}((c, i), x) = A(c, i), \forall ((c, i), x)$ and, thus, $\tilde{A} = \cup_{((c,i),x) \in \tilde{S}} \tilde{A}((c, i), x) = \cup_{(c,i) \in S} A(c, i) = A$. In a similar way, the reward set \tilde{R} is build from R . For $\tilde{s} \in \tilde{S}, a \in A$, the reward $\tilde{r}_n(\tilde{s}, a) \in \tilde{R}$ is $\tilde{r}_n(((c, i), x), a) = aF_i$ for $c, i > 0$ and 0 otherwise. Thus, $\tilde{R} = R$.

As well, $\tilde{P} = P$, as the transition probabilities \tilde{P} are determined by P . We have $\tilde{p} \in \tilde{P}$ and, with states $((c, i), x), ((c - a, j), x - aF_i) \in \tilde{S}$ and $a \in A$, the transition probability is given by $\tilde{p}_n(((c - a, j), x - aF_i) | ((c, i), x), a) = p_{n,j}$ for $n = N, N - 1, \dots, 0$ and else 0.

We are interested in the probability that our obtained total revenue does not attain a target level x . Let the set of deterministic Markovian policies be $\tilde{\Pi}$ and let the random variable for the cumulative gained reward, applying policy $\tilde{\pi} \in \tilde{\Pi}$ beginning with capacity c , request i , remaining time steps N , and target x , be $\rho_N^{\tilde{\pi}}((c, i), x) = \sum_{n=0}^N \tilde{r}_n$. For the policy $\tilde{\pi}$, the target-percentile risk measure is defined as

$$V_N^{\tilde{\pi}}((c, i), x) := \mathcal{P}(\rho_N^{\tilde{\pi}}((c, i), x) \leq x), \tag{1}$$

where \mathcal{P} denotes a probability. The time consistency property of the target-percentile risk measure can be shown as demonstrated by Boda and Filar (2006).

Thus, we are looking now for an optimal policy $\tilde{\pi}^*$ for each objective function $V_N^{\tilde{\pi}}((c, i), x)$ that minimises the risk of failing target level x :

$$\tilde{\pi}^* \in \operatorname{argmin}_{\tilde{\pi} \in \tilde{\Pi}} \{V_N^{\tilde{\pi}}((c, i), x) \mid ((c, i), x) \in \tilde{S}\}. \quad (2)$$

The associated percentile (minimum risk level for x) is denoted $V_N^{\tilde{\pi}^*}$.

Following Wu and Lin (1999) and Boda and Filar (2006), we can derive the dynamic programming equations for computation of the minimum percentile $V_N^{\tilde{\pi}^*}$ (see Appendix). We attain the following equations for $x \in \mathbb{R}$, $\forall (c, i) \in S$:

$$V_0^{\tilde{\pi}^*}((c, i), x) = \begin{cases} 1 & x > 0 \\ 0 & \text{otherwise,} \end{cases}$$

$$V_n^{\tilde{\pi}^*}((c, i), x) = \min_{a \in A} \left\{ \sum_{j=0}^k p_{n-1,j} V_{n-1}^{\tilde{\pi}^*}((c-a, j), x-aF_i) \right\},$$

$$n \in \{1 \dots N\}. \quad (3)$$

In time $n=0$, the initial probabilities are one for a target $x > 0$ (as there is no remaining time for earning any value) and 0 for a target $x \leq 0$ (as this will definitely be met because our initial revenue is zero). Note that the final percentile of all time periods is determined by $V_{N+1}^{\tilde{\pi}^*}(c, i, x)$; in $V_N^{\tilde{\pi}^*}(c, i, x)$, we already know the requested class i at time N .

The optimal policy $\tilde{\pi}^*$ can be computed from the minimum percentile $V_N^{\tilde{\pi}^*}$ by Equation (2) for a given target level x . It should be pointed out that an optimal policy describes one way to obtain the target level, but several optimal policies might exist. Therefore, if more than one decision rule can be chosen in a certain state in order to achieve the minimum percentile we select the decision rule which contributes most to the revenue. In particular, we prefer to accept a request if the probabilities of both possible decisions $a \in \{0, 1\}$ are equal when determining the minimum in Equation (3), $\sum_{j=0}^k p_{n-1,j} V_{n-1}^{\tilde{\pi}^*}((c, j), x) = \sum_{j=0}^k p_{n-1,j} V_{n-1}^{\tilde{\pi}^*}((c-1, j), x-F_i)$ and the risk-neutral solution would accept the request, too. In this manner, we achieve the same probabilities regarding the target but with the policy which yields the greater expected revenue.

Furthermore, if we have in some state achieved the target level the following states can be arbitrarily chosen. In practice, the policy for the ongoing states should be optimised then under another criterion, such as the expected revenue. Moreover, if the target can never be obtained in the given setting, all policies are equally improper and no optimal target-percentile policy exists (technically, all policies are optimal but none is proper). In both cases, we apply the risk-neutral policy which maximises the expected revenue throughout this paper if not otherwise stated.

For efficient implementation, we apply an usual transformation of the dynamic programming formulation of Equation (3). Introducing the operator $T_n(c, x) := \sum_{i=0}^k p_{n,i} V_n(c, i, x)$ helps reducing the state space by variables representing the fare class of an arrival. Defining $W_n(c, x) := T_n(c, x) V_n(c, i, x)$, we transform Equation (3), as follows, for $x \in \mathcal{R}$, $c \in \{0, \dots, C\}$:

$$W_0^{\tilde{\pi}^*}(c, x) = \begin{cases} 1 & x > 0 \\ 0 & \text{otherwise,} \end{cases}$$

$$W_n^{\tilde{\pi}^*}(c, x) = T_n(c, x) V_n^{\tilde{\pi}^*}(c, i, x)$$

$$= \sum_{i=0}^k p_{n,i} \min_{a \in A} \left\{ W_{n-1}^{\tilde{\pi}^*}(c-a, x-aF_i) \right\}. \quad (4)$$

The computation of all possible cumulative rewards given by the variable x could be reduced if done on a suitable grid for larger problems as described in the works of Wu and Lin (1999) and Boda *et al* (2004). In this paper, we do not apply the grid reduction.

Example. In order to illustrate the method, we give a stylised example. Consider only two classes with fares $F_1 = 200$; $F_2 = 100$, two remaining time periods $N=2$, one seat left $C=1$, and the probabilities for arrivals $p_{1,1} = 0.10$, $p_{1,2} = 0.15$, $p_{2,1} = p_{2,2} = 0.20$. Thus, for example, the probability of a request of fare 2 in period 1 before departure is 15%. We have a few scenarios in this setting: if a request for a distinct fare class comes in period 2 before departure, we can accept it or reject this fare class and then wait for possible arrivals in the last period and, if they appear, accept. It is easy to see that the policy which always accepts (expected revenue of 81) is better off when compared with others. However, consider that now we want the best policy for a target value of 200. The expected revenue maximising policy fails that target with probability of 0.74. A better choice for this target would be only acceptance of the highest fare class, a policy which fails only with a likelihood of 0.72.

4. Numerical simulation and results

In their introductory paper about dynamic capacity control, Lee and Hersh (1993) used an example which also served for illustration in the recent papers of Barz and Waldmann (2007), Huang and Chang (2011) and Koenig and Meissner (2015). Thus, we can also demonstrate the proposed target-percentile policy in the same exemplary setup.

4.1. Exemplary simulation setup

There are $N=30$ time periods before departure, and the initial number of seats is $C=10$. The four fare classes are $F_1 = 200$, $F_2 = 150$, $F_3 = 120$, $F_4 = 80$. The probabilities for a request of a fare class in a given time period are shown in Table 1.

In order to see how the target-percentile policy works, we conducted an experiment with 1000 sample runs. Random arrivals were simulated in a Monte Carlo manner using the values of Table 1. When compared with other proposed policies, the same sample paths (random arrivals) were used, of course.

A single simulation run is initialised with values for remaining seats, time periods before departure, and a policy. The policy contains for each state the acceptable fare classes. The state is described by remaining time periods, remaining seats, and a remaining target value. The simulation then continues to loop over the time periods until the departure time zero is reached. Inside the loop, a random generator simulates requests for fare classes which are accepted if the current policy allows acceptance of the class or else rejected. An update of the state is as follows: time periods are always decremented by one, seats are decremented only if a fare is accepted, and the target value is decremented by the gained fare price.

Table 1 Fares and request probabilities for fare class i and time period n

i	F_i	$1 \leq n \leq 4$	$5 \leq n \leq 11$	$12 \leq n \leq 18$	$19 \leq n \leq 25$	$26 \leq n \leq 30$
1	200	0.15	0.14	0.10	0.06	0.08
2	150	0.15	0.14	0.10	0.06	0.08
3	120	0	0.16	0.10	0.14	0.14
4	80	0	0.16	0.10	0.14	0.14

Policy illustration. Figure 1 visualises the policy $\tilde{\pi}^*$ for the described example. We see slices through a three-dimensional matrix which displays the index of the maximum allowed fare class for each state (c, x) in time n with initial target of 1200. In order to use the policy, we start in the state $(10, 1200)$ with 30 time periods to go. This is the top corner on the right hand side of the presented box. The state at this position in the matrix gives the maximum allowed fare class, which lets one decide how to act at this point in time before departure. Only fare classes with higher or equal price than the associated class shown are accepted. As time marches on, one moves always one step further along the time dimension to departure time zero; this is parallel to the south-west direction in the figure. The policy decides now which way to move in both other dimensions. An acceptance of a request causes a move downwards along the dimension of the capacity, orthogonally downwards in the matrix. Finally, the price of an accepted fare dictates how to move in the target direction, along the north-west direction in the figure. Thus, considering the figure, the simulation will generate random trajectories from the top corner on the right hand side to the bottom corner on the left hand side. Of course, the end of each trajectory will often be different because of the random realisations but it has to end with coordinate $n = 0$.

We illustrate the effect of changing a target level on the policy in Figure 2. The figure shows for two different target levels (1200, 1400) the corresponding policies when revenue has not yet been gained. The effect of increasing the target can




Figure 1 Target level policy: Indices of maximum allowed fare class of policy $\tilde{\pi}^*$ are shown for each state (c, x) with remaining capacity $0 \leq c \leq 10$, revenue target $0 \leq x \leq 1200$ in time $n, 1 \leq n \leq 30$. The indices of the four fare classes are visualised by greyscale where brightness increases with fare class. For example, class 1 is darkest. Only fare classes of higher or equal price than indicated by the greyscale are accepted in a state. Note, states which offer no optimal solution because the revenue is not obtainable, given the remaining capacity and time, are transparent.




Figure 2 Illustration of the effect by changing the target revenue on the policy: The figures show the indices of the maximum allowed fare class of the policies for a capacity between 1 and 10 seats, a remaining time between 1 and 30 periods and when revenue has not yet been collected. Only fare classes up to the indicated level are accepted in a state. States which do not allow a policy to attain the target are shown in black colour.

- (a) Target $x_N = 1200$.
- (b) Target $x_N = 1400$.

be observed by the right hand side of the matrices which show the indices of the maximum allowed fare class for each state. For example, a capacity of at least six seats is required for a target of 1200, but a capacity of at least seven seats is required for a target of 1400. Only the highest fare class is accepted when only six seats are available for the target of 1200, respectively, when only seven seats are available for the target of 1400. Figure 2 illustrates the pure target policies for 1200 and 1400. As already mentioned, we apply the risk-neutral policy in states which do not allow to achieve a target; this is not shown in this figure.

Evaluation. As the proposed policy optimises the target-percentile, we start our evaluation with different (obtainable) target revenues, comparing the theoretical and the simulation results. As mentioned in Section 3.2, there are scenarios when target revenue is achieved but time is remaining and one or more seats are left. We present the average of remaining time and seats for such cases as well. Further, the averaged revenue is computed by switching to the risk-neutral policy or the first-come-first-serve (FCFS) policy when the target has been achieved. Table 2 shows the results for seven different targets. The average of failed cases in the simulation is plausible within numerical errors to the theoretical percentile, validating that the policy does as expected.

The expected revenue of the risk-neutral policy for the analysed problem is 1407.2. Looking at the results of Table 2, we see that a policy which aims towards a lower target revenue than the expected value accepts an upcoming request early in time. Decisions are made soon and not post-poned to later periods. This effect is easily observable as remaining time and seats decrease, while the target is increasing. Policies with lower targets have a greater probability for reaching the target. It can be more easily obtained by accepting requests early, thus leaving more time for balancing against having no profitable requests in the next time periods. For the very low target of 800, the target policy is similar to the FCFS policy and every request is taken early in order to achieve the low target. For the very high target of 1900, the target policy ‘speculates’ for unlikely

combinations of requests of high fares and leaves with empty seats.

The effect between switching to the risk-neutral or FCFS policy for the remaining time after achieving the target can be observed for the revenue and standard deviation. Of course, there is no impact on failed target, remaining time and seats. With decreasing remaining time (or increasing target), the difference between the revenue and standard deviation of using the risk-neutral or FCFS policy for the remaining time diminishes. The average revenues of the target policies are in each case lower than that of the risk-neutral policy but greater than the FCFS policy. The standard deviation of these revenues grows with an increasing target, although when compared with the risk-neutral and FCFS cases, their policies less often fail the targets. This can be explained by comparing the distribution histograms of the revenues of the policies. In the following, we apply the risk-neutral policy when a target is reached. Figure 3 shows the distribution histograms of 1000 simulation runs of three policies: one with low target 1200, one with high target 1400, and a risk-neutral one maximising expected revenue.

The distribution associated with the low target has its peak above its target value 1200 and a slight negative skew. It has only small frequencies for values lower than 1200 but also for values higher than 1500, as its standard deviation from Table 2 also emphasises. It has a peak at 1300. The risk-neutral solution shows a negative skewed distribution with a peak at 1500 with a long tail to very low values, though some high revenues at 1800. Compared with the policy with target 1200, its revenues are more often below 1200; however, given the revenue is greater 1200, it will be better off. Its risk of falling below 1200 remains higher than the risk of the low target policy.

The distribution of the policy with high target 1400 has a strong negative skew with a long tail to low values, too. The peak of the distribution is at 1400. Compared with the risk-neutral counterpart, this policy shifts frequency from 1300 to 1400 revenue. The target is achieved mainly at the expense of 1300 revenue and greater than 1500 revenue. Further, it shows also higher frequencies for low revenue than both other policies. Hence, if it fails the target, there is a greater risk of obtaining only low revenue.

Table 2 Results of policy simulation for different target values. The probabilities and averages for failing to achieve a target are given (lower means better) and also, averages of remaining time and seats if target could be achieved

<i>Target value</i>	1900	1600	1500	1400	1300	1200	800
$V_{\bar{x}}^*$	0.993	0.740	0.528	0.336	0.183	0.088	0.004
<i>Simulation of target value policies</i>							
Failed target	0.998	0.749	0.510	0.354	0.189	0.117	0.006
Remaining time	0	0.04	0.22	0.82	1.93	3.87	13.92
Remaining seats	0	0.01	0.04	0.11	0.23	0.49	3.16
<i>Using risk-neutral policy for remaining time</i>							
Revenue	1006.2	1314.8	1362.5	1367.4	1345.8	1319.9	1354.2
Standard deviation	291.15	276.97	242.76	191.56	161.23	160.81	177.26
<i>Using FCFS policy for remaining time</i>							
Revenue	1006.2	1314.6	1361.6	1365.4	1338.3	1307.0	1290.4
Standard deviation	291.15	276.57	241.59	188.48	151.86	146.46	155.88
<i>Simulation of risk-neutral policy</i>							
Failed target	1.000	0.821	0.621	0.420	0.271	0.158	0.014
Revenue				← 1402.2 →			
Standard deviation				← 210.36 →			
<i>Simulation of FCFS policy</i>							
Failed target	1000	0.982	0.911	0.742	0.504	0.270	0.006
Revenue				← 1290.4 →			
Standard deviation				← 155.88 →			




Figure 3 Histogram of the distributions of revenue in the simulation using policies with targets 1200 (left bars), 1400 (middle bars) and the risk-neutral policy (right bars).

The histogram demonstrates that the policy with low target aims at a lower average revenue and smaller variance, but the policy with a higher target, near to the expected revenue of the risk-neutral solution, does not.

In order to evaluate the performance of target revenue policies in more detail, we compare them with the risk-

sensitive policies derived from expected utility theory, as in Barz and Waldmann (2007). We select the latter policies for comparison as they result from optimising the dynamic capacity control model using an exponential utility and no heuristics. Referring to the recent works of Huang and Chang (2011) and Koenig and Meissner (2015), we view the mean, standard deviation, and CV@R of the policies. The CV@R is a measure for the expected revenue given the revenue is below a certain quantile specified by a confidence level α ; it is the expected value in the α percent of worst cases.

Table 3 compares both types of risk-sensitive policies. Beyond the mean, standard deviation, and CV@R with confidence level 5%, the observed relative frequency of failing the 1000 target is given. We see that the target policy for 1000 has the least risk of failing it. However, it is also observable that the target policies only limit the risk of failing their certain target and do not provide more preferable results in terms of the other measures. The expected utility based policies have higher average revenues than most target policies. If the target policies aim at a level greater than 1200, the CV@R drops down with further increasing the level. The CV@R of the policies employing an exponential utility function decreases with decreasing level of risk aversion. The standard deviation decreases with higher target and higher level of risk sensitivity for both types of policies, with exception of the 1400 target. Discussed already by Figure 3, the CV@R results also show that the target policies do not limit the risk of obtaining only low revenues in the worst cases. Further, it is interesting that the policies aimed at targets

Table 3 Comparison between two risk-sensitive policies: target-percentile optimising and exponential utility function optimising policies (the risk aversion increases in conjunction with γ). CV@R is for $\alpha = 5\%$

Policy	Mean	Standard deviation	CV@R	Frequency ($\rho < 1000$)
Target 800	1354.2	177.3	941.2	0.028
Target 900	1349.6	171.8	960.4	0.029
Target 1000	1331.6	170.3	959.4	0.020
Target 1100	1322.0	165.5	951.8	0.031
Target 1200	1319.9	160.8	927.0	0.029
Target 1300	1345.8	161.2	910.2	0.041
Target 1400	1367.4	191.6	874.0	0.050
Utility $\gamma = 0.010$	1359.5	162.3	953.2	0.027
Utility $\gamma = 0.005$	1383.8	176.3	939.6	0.029
Utility $\gamma = 0.001$	1402.3	206.8	886.6	0.050

different from 1000 do not guarantee good performance regarding the 1000 target.

This effect becomes more observable in the distribution histogram of the 1000 revenue target policy and the expected utility policy with high risk aversion $\gamma = 0.01$, as shown in Figure 4. The target policy has a lower average revenue, a higher 5% CV@R, and a higher standard deviation than the exponential utility policy, and it achieves at least a revenue of 1000 in more cases. The frequencies for the revenues 800 and 900 are lower for the target policy than for the exponential utility policy. The target policy has higher frequencies for revenues between 1000 and 1200 and between 1700 and 1800. It has lower frequencies between 1300 and 1600 than its counterpart. This explains the lower mean revenue.

Figures 3 and 4 show that the target policies dent the distribution slightly below the target. Thereby, the whole distribution, for values lower and greater than the target, is influenced. Frequencies below this dent may increase as frequencies for the target do. In particular, distribution lower than the target need not be modified in a favourable manner regarding the lowest revenues, that is to say the worst cases.

The results of Table 3 show, that decisions makers should choose a policy according their prioritisation of measures. For example, a risk-averse policy is appropriate for decision makers if their business could be more negatively impacted by a (few) worst case scenarios than forgoing revenue in average.

Further numerical experiments. We did further numerical experiments beyond the previous illustrative one. In order to investigate the target policy, we show five more scenarios which differed with respect to their load factor. The load factor is given by $\lambda = 1/C \sum_{n=0}^N \sum_{j=1}^k p_{jn}$ and gives information about demand in relation to capacity. The previous example had a load factor of $\lambda = 1.32$.

We changed only the request probabilities of the previous example and hold the other parameters fixed to get further




Figure 4 Histogram of the distributions of revenue in the simulation with policy with target 1000 (left bars) and policy using exponential utility function with $\gamma = 0.01$ (right bars).

scenarios. To this end, we built the further scenarios by choosing random request probabilities which yielded different load factors. Then we simulated 1000 sample runs with each scenario. Table 4 show results of the risk-neutral policy and of the target policy which were applied to the scenarios. We selected those targets for each policy which were 115, 100 and 85% times the expected revenue of the risk-neutral policy. As Table 4 shows, the target policies achieved the desired target more often than the risk-neutral policies in the numerical simulations. That advantage of the target policies increased along with the increasing load factor.

5. Conclusions

A risk-averse policy minimising the failure of a previously defined, certain revenue target has been proposed for a revenue management problem, namely the dynamic capacity control setting. This policy is derived by extending the state space of the Markov decision process formulation of the problem. We have discussed aspects for implementing the policy numerically. In numerical experiments, we have analysed the proposed policy and evaluated against risk-neutral and another risk-sensitive policies. We have compared the mean, standard deviation, and conditional-value-of-risk of those policies. The optimal policy for a given target revenue focuses on minimising the likelihood of the failing of this certain target but does not compensate for other risk measures.

The analysis of the revenue distributions of the target revenue aimed policies in numerical experiments disclose how important correct understanding of such a policy is when applied. The decision maker must be aware of its limitations; in particular, that it is the policy with lowest probability of failing the target, but the probability of worst outcomes are not eliminated.

Table 4 Results of numerical simulation with scenarios which differed by their load factors λ . Revenues are averaged over 1000 sample runs. The given differences are the observed relative frequencies of failed target instances of the risk-neutral policy minus the observed relative frequencies of failed target instances of the target policy. The differences show how more often the target level policies achieved the target but the risk-neutral ones did not. For example, in the example of the last row of the table, the target policy did failed in 1.4% of all sample runs and the risk-neutral one in 7.1%

Target	V_n^*	Failed target	Difference failed versus risk neutral	Revenue
$\lambda = 0.89$, expected revenue = 1157.2, simulation revenue = 1155.9				
1340	0.697	0.699	0.001	1145.0
1160	0.453	0.449	0.022	1151.5
990	0.255	0.258	0.002	1156.6
$\lambda = 0.99$, expected revenue = 1280.1, simulation revenue = 1277.8				
1480	0.750	0.750	0.038	1251.6
1290	0.453	0.470	0.070	1259.2
1090	0.210	0.212	0.013	1271.6
$\lambda = 1.11$, expected revenue = 1356.3, simulation revenue = 1358.7				
1560	0.779	0.776	0.072	1320.2
1360	0.398	0.387	0.046	1345.1
1160	0.131	0.119	0.039	1335.8
$\lambda = 1.48$, expected revenue = 1482.0, simulation revenue = 1484.3				
1710	0.828	0.821	0.073	1391.4
1490	0.386	0.382	0.070	1446.2
1260	0.070	0.075	0.037	1389.6
$\lambda = 1.76$, expected revenue = 1525.5, simulation revenue = 1522.7				
1760	0.931	0.934	0.034	1269.2
1530	0.333	0.320	0.160	1477.9
1300	0.012	0.014	0.057	1437.4

However, using a low target revenue supports limiting such risk.

The presented approach can be further developed in order to achieve a policy which optimises value-at-risk as proposed by Boda and Filar (2006). Furthermore, it also offers the basis for the development of investigating policies balancing out mean revenue *versus* target achievement.

References

- Barz C and Waldmann K-H (2007). Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research* **65**(3): 565–579.
- Beloba PP (1989). Application of a probabilistic decision model to airline seat inventory control. *Operations Research* **37**(2): 183–197.
- Bitran G and Caldentey R (2003). An overview of pricing models for revenue management. *Manufacturing & Service Operations Management* **5**(3): 203–229.
- Boda K and Filar JA (2006). Time consistent dynamic risk measures. *Mathematical Methods of Operations Research* **63**(1): 169–186.

- Boda K, Filar JA, Lin Y and Spanjers L (2004). Stochastic target hitting time and the problem of early retirement. *IEEE Transactions on Automatic Control* **49**(3): 409–419.
- Bouakiz M and Kebir Y (1995). Target-level criterion in markov decision processes. *Journal of Optimization Theory and Applications* **86**(1): 1–15.
- Feng Y and Xiao B (1999). Maximizing revenues of perishable assets with a risk factor. *Operations Research* **47**(2): 337–341.
- Feng Y and Xiao B (2008). A risk-sensitive model for managing perishable products. *Operations Research* **56**(5): 1305–1311.
- Gönsch J and Hassler M (2014). Optimizing the conditional value-at-risk in revenue management. *Review of Management Science* **8**(4): 495–521.
- Huang K and Chang K (2011). A model for airline seat control considering revenue uncertainty and risk. *Journal of Revenue and Pricing Management* **10**(2): 161–171.
- Koenig M and Meissner J (2010). List pricing versus dynamic pricing: Impact on the revenue risk. *European Journal of Operational Research* **204**(3): 505–512.
- Koenig M and Meissner J (2015). Risk management policies for dynamic capacity control. *Computers and Operations Research* **59**: 104–118.
- Lancaster J (2003). The financial risk of airline revenue management. *Journal of Revenue and Pricing Management* **2**(2): 158–165.
- Lautenbacher CJ and Stidham SJ (1999). The underlying markov decision process in the single-leg airline yield-management problem. *Transportation Science* **33**(2): 136–146.
- Lee T and Hersh M (1993). A model for dynamic airline seat inventory controls with multiple seat bookings. *Transportation Science* **27**(3): 252–265.
- Levin Y, McGill J and Nediak M (2008). Risk in revenue management and dynamic pricing. *Operations Research* **56**(2): 326–343.
- Maglaras C and Meissner J (2006). Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing & Service Operations Management* **8**(2): 136–148.
- Shapiro A (2009). On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters* **37**(3): 143–147.
- Talluri K and van Ryzin G (2005). *The Theory and Practice of Revenue Management*. Springer: New York.
- Weatherford L (2004). EMSR versus EMSU: Revenue or utility? *Journal of Revenue and Pricing Management* **3**(3): 277–284.
- Weatherford L and Beloba PP (2002). Revenue impacts of fare input and demand forecast accuracy in airline yield management. *Journal of the Operational Society* **53**(8): 811–821.
- White DJ (1988). Mean, variance, and probabilistic criteria in finite markov decision processes: A review. *Journal of Optimization Theory and Applications* **56**(1): 1–29.
- Wu C and Lin Y (1999). Minimizing risk models in markov decision processes with policies depending on target values. *Journal of Mathematical Analysis and Applications* **231**(1): 47–67.
- Xiong H, Xie J and Deng X (2011). Risk-averse decision making in overbooking problem. *Journal of the Operational Society* **62**(9): 1655–1655.

Appendix

The following theorems are based on the work of Bouakiz and Kebir (1995). The theorems show that our dynamic programming equations find the optimal policy.

Theorem 1 For each $(c, i) \in S$ and $n \geq 0$, $V_n^*((c, i), \cdot)$ is a distribution function.

$V_0^{\pi^*}((c, i), \cdot)$ is a distribution function. Assume that $V_{n-1}^{\pi^*}((c, i), \cdot)$ is a distribution function. As by definition

$$V_n^{\pi^*}((c, i), x) = \min_{a \in A} \left\{ \sum_{j=0}^k p_{n-1,j} V_{n-1}^{\pi^*}((c-a, j), x - aF_i) \right\}.$$

The sets A, S, R are all finite. The finite convex combination and minimum of distribution functions are distribution functions, too. Thus, $V_n^{\pi^*}((c, i), \cdot)$ is a distribution function.

Theorem 2 For each $n \geq 1$, the sequence $\{V_n^{\pi^*}\}$ satisfies the relations

$$V_n^{\pi^*}((c, i), x) = \begin{cases} 1, & \text{if } x > \rho_N^{\pi^*}(c, i) \\ \min_{a \in A} \left\{ \sum_{j=0}^k p_{n-1,j} V_{n-1}^{\pi^*}((c-a, j), x - aF_i) \right\}, & \text{if } x \leq \rho_N^{\pi^*}(c, i). \end{cases}$$

We define the following operators for convenience. Let a be an arbitrary action and δ_n an arbitrary decision rule at time n . The operators M^a, M^{δ_n} and M are as follows:

$$M^a g_n((c, i), x) = \sum_{j=0}^k p_{n-1,j} g_{n-1}((c-a, j), x - aF_i),$$

$$M^{\delta_n} g_n((c, i), x) = M^{\delta_n((c,i),x)} g_n((c, i), x),$$

$$M g_n((c, i), x) = \min_{a \in A} \{M^a g_n((c, i), x)\}.$$

Let $\pi = (\delta_{n-1}, \dots, \delta_1)$ be a policy and δ_n be a decision rule. Let policy $\gamma = (\delta_n, \delta_{n-1}, \dots, \delta_1)$, then $V_n^\gamma((c, i), x) = M^{\delta_n} V_{n-1}^\pi((c, i), x)$ by definition. With $a = \delta_n((c, i), x)$, it is

$$V_n^\gamma((c, i), x) = M^a V_{n-1}^\pi((c, i), x).$$

It follows that

$$V_n^\gamma((c, i), x) \geq M V_{n-1}^\pi((c, i), x) \geq M V_{n-1}^{\pi^*}((c, i), x),$$

and the minimum over γ yields

$$V_n^{\pi^*}((c, i), x) \geq M V_{n-1}^{\pi^*}((c, i), x).$$

With γ, a and $V_n^\gamma((c, i), x) = M^a V_{n-1}^\pi((c, i), x)$ as before, it is

$$V_n^{\pi^*}((c, i), x) \leq M^a V_{n-1}^{\pi^*}((c, i), x), \forall \pi.$$

Let $\epsilon > 0$ be arbitrary and λ be a policy which depends on $(c, i), x, \epsilon$ such that

$$V_{n-1}^\lambda((c, i), x) \leq V_{n-1}^{\pi^*}((c, i), x) + \epsilon.$$

As π is arbitrary in the above formula,

$$V_n^{\pi^*}((c, i), x) \leq M^a V_{n-1}^\lambda((c, i), x)$$

$$\leq M^a [V_{n-1}^{\pi^*}((c, i), x) + \epsilon] \leq M^a V_{n-1}^{\pi^*}((c, i), x) + \epsilon.$$

Now, since δ_n is an arbitrary decision rule, $a = \delta_n((c, i), x)$, it follows

$$V_n^{\pi^*}((c, i), x) \leq M V_{n-1}^{\pi^*}((c, i), x) + \epsilon,$$

and $\epsilon > 0$ is arbitrary,

$$V_n^{\pi^*}((c, i), x) \leq M V_{n-1}^{\pi^*}((c, i), x).$$

Hence,

$$V_n^{\pi^*}((c, i), x) = M V_{n-1}^{\pi^*}((c, i), x).$$

Note, the sequence $\{V_n^{\pi^*}\}$ is monotonically decreasing. It is $V_1^{\pi^*} \leq V_0^{\pi^*}$. Hence, assuming that $V_n^{\pi^*} \leq V_{n-1}^{\pi^*}$, we have $V_{n+1}^{\pi^*} = M V_n^{\pi^*} \leq M V_{n-1}^{\pi^*} = V_n^{\pi^*}$.

Received 7 October 2014;
accepted 8 July 2015 after one revision



This work is licensed under a Creative Commons Attribution 3.0 Unported License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/3.0/>