Int. J. Production Economics 166 (2015) 11–19

Contents lists available at ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

Value-at-risk optimal policies for revenue management problems

Matthias Koenig [a,*], Joern Meissner [b]

[a] Department of Management Science, Lancaster University Management School, United Kingdom
[b] Kuehne Logistics University, Hamburg, Germany

ARTICLE INFO

Article history:
Received 12 June 2011
Accepted 25 March 2015
Available online 11 April 2015

Keywords:
Capacity control
Revenue management
Risk
Value-at-risk

ABSTRACT

Consider a single-leg dynamic revenue management problem with fare classes controlled by capacity in a risk-averse setting. The revenue management strategy aims at limiting the down-side risk, and in particular, value-at-risk. A value-at-risk optimised policy offers an advantage when considering applications which do not allow for a large number of reiterations. They allow for specifying a confidence level regarding undesired scenarios.

We introduce a computational method for determining policies which optimises the value-at-risk for a given confidence level. This is achieved by computing dynamic programming solutions for a set of target revenue values and combining the solutions in order to attain the requested multi-stage risk-averse policy. We reduce the state space used in the dynamic programming in order to provide a solution which is feasible and has less computational requirements. Numerical examples and comparison with other risk-sensitive approaches are discussed.

1. Introduction

Revenue management deals with controlling a revenue stream resulting from selling products using a fixed and perishable resource. The industries which use revenue management are manifold. The most popular representatives are airlines, hotels, rental cars, and advertising. But revenue management is also common in event management, ferry lines, retailing or healthcare, to name a few. Talluri and van Ryzin (2005) and Chiang et al. (2007) provide a comprehensive overview of revenue management.

The firm sells multiple products, each consuming a fixed resource with a limited capacity. In this setting, we consider quantity-based revenue management in which a company offers all or just a subset of all products at each point in time. There is a finite time horizon for selling the products, as at the end of the horizon, the salvage value of the resource is zero.

The most common settings use the assumption of a risk-neutral objective. Thus, the policy of the firm is the maximisation of the expected value of its revenue. Often, such a risk-neutral objective is sufficient. As in most applications, such as daily operating ferry lines, this policy is repetitively used. By the law of large numbers, using the expected value as the objective function is then appropriate.

Nevertheless, risk neutrality may not be adequate for other industries, such as event management, that do not support a large number of repetitions of a policy. Several scenarios are known that argue for the considerations of risk-sensitive or risk-averse policies.

Levin et al. (2008) emphasise that, in particular, an event promoter has a high risk, as the promoter cannot count on a large number of reiterations of events. The promoter faces high fixed costs and predominantly has to recover them in order to avoid a possible high loss. Financial and also strategic reasons might not allow running into negative cash, because operational mobility might suffer.

Both Bitran and Caldentey (2003) and Weatherford (2004) provide further examples that risk-neutral considerations are not applied for every real scenario. They report that airline analysts show some natural risk-averse behaviours, and they overrule their revenue management system in situations when the system recommends waiting for high-fare passengers, instead accepting low-fare passengers a few days before flight departure.

That risk-neutral and risk-sensitive policies make a difference that is shown in several recent papers. Barz and Waldmann (2007), Huang and Chang (2011), and Koenig and Meissner (2013, 2015) analyse both types of policies using the same underlying model that is used in this paper. All four approaches analyse the effects of applying different kinds of risk-sensitive polices, assuming various levels of risk aversion for a decision maker. However, none of these approaches computes an optimal policy for the common risk measures, such as standard deviation, value-at-risk, or conditional-value-at-risk. However, simulations can be run to determine their values for a given policy.

In this paper, we propose a method which computes a value-at-risk optimal policy. The value-at-risk ($V@R$) is a common risk measure

* Corresponding author at: Bielefeld University of Applied Sciences, Germany.
E-mail address: matthias.koenig@fh-bielefeld.de (M. Koenig).

http://dx.doi.org/10.1016/j.ijpe.2015.03.027
0925-5273/

often used in finance (cf. Jorion, 2006). It measures down-side risk and is determined for a given probability level. With regard to $V@R$, this probability level is often referred to as a confidence level. In our context, the $V@R$ is the lowest revenue which exceeds the confidence level, which is often set at 5 or 10%. Basically, it is a quantile of the revenue distribution determined by the given confidence level. However, $V@R$ misses the subadditivity property (cf. Rockafellar and Uryasev, 2000) and, thus, conditional-value-at-risk has become very popular as risk measure, e.g. Yau et al. (2011) use it for financial and operational decisions in the electricity sector.

Nevertheless, in order to find a $V@R$ optimal policy, we take advantage of the computation of target level optimal policies as proposed by Koenig and Meissner (2013). The target level optimal policy can be computed for a certain target and gives information about the probability of not achieving this target. This probability is minimised to find the best policy. It defines a confidence level for a fixed target, which is the corresponding $V@R$. Hence, our task is similar to computing a target level optimised policy, but we optimise the threshold value instead of the percentile. We compute $V@R$ optimal policies and their associated confidence levels. We determine then the policy of the desired confidence level by evaluating the confidence levels of these policies. We describe in this paper how that can be accomplished in an efficient manner.

The advantage of using $V@R$ as a parameter to be optimised is that it is a well-known risk measure, and it is easily interpreted by practitioners. A desired confidence level is specified, and the $V@R$ is returned in the monetary unit of the revenue. Other risk-sensitive approaches often require an interpretation of an uncommon parameter to adjust the desired level of risk preference. $V@R$ is well established and used by risk analysts and decision makers as a standard tool not only for financial investments. The risk of a strategy pursued by a decision maker can be assessed by a clear definable risk exposure. This enables risk assessments and planning on an organisational level. Managers can choose their confidence level and communicate it to upper management and investors as well.

Further, a decision maker can define the confidence level to be used for a range of problems although the problems might differ in their settings. This is a great benefit of the $V@R$ approach when compared with the target level approach which might require different target values for each problem setting.

The contribution of this paper is a novel approach in order to assess risk in a revenue management setting. Our approach computes efficiently a value-at-risk optimal policy. To this purpose, we introduce an innovative method in order to reduce the state space of the method which computes a target level optimal policy. We present a simulation study which highlights that our state space reduction still yields high accuracy for the $V@R$ computation even with a significant decreased number of states. In this way, we deliver also a solution which is feasible and has less computational requirements.

The paper is structured as follows. This introduction is followed by a brief overview of related work dealing with revenue management models incorporating risk in Section 2. In Section 3, we continue with the description of the revenue model, which builds our basic position. We describe the target level approach and how we use it to efficiently obtain a $V@R$ optimal policy. We discuss different strategies useful for numerical approximation of such a policy. Section 4 gives a detailed overview of the numerical results and studies the effect of numerical approximation methods. Finally, we conclude this paper in Section 5.

2. Related work

As a starting point for our analysis we use the basic model by Lee and Hersh (1993). They introduce the dynamic capacity control model in a risk-neutral setting. Lautenbacher and Stidham (1999) take this model further and derive a corresponding Markov decision process. This description as a Markov decision process is advantageous for model extensions.

First risk considerations in revenue management models are proposed by Feng and Xiao (1999). Their model considers risk in terms of variance of sales due to changes of prices. To this end, a penalty function reflecting this variance is incorporated in the objective function of the model. Further, Feng and Xiao (2008) integrate expected utility theory into revenue management models in order to support risk-sensitive decisions.

Expected utility theory as a tool for risk consideration is recommended by Weatherford (2004), as well. From a practitioner's perspective, he criticises risk-neutral revenue management, in particular, the expected marginal seat revenue (EMSR) heuristic by Beloba (1989), and endorses risk-averse models.

Barz and Waldmann (2007) base their risk-sensitive model on the Markov decision process of the dynamic capacity model and expected utility theory. They integrate an exponential utility function as the objective function into the Markov decision model. The exponential utility function allows the use of different levels of risk sensitivity. Barz (2007) points out the use of a utility function with an aspiration level in the same setting but does not discuss the computation of an optimal policy for this utility function. Maximising expected utility using an aspiration level states the same problem as done by the target level objective which is discussed in this paper.

Further, Gönsch and Hassler (2014) deal with finding an optimal conditional-value-at-risk policy and derive an heuristic in which a solution of a continues knapsack problem is required in each state of their value function.

Another way of employing expected utility theory in a revenue management context is proposed by Lim and Shanthikumar (2007). They analyse robust and risk-sensitive control with an exponential utility function for dynamic pricing.

Lai and Ng (2005) formulate a robust optimisation model for revenue management in the hotel industry. Their model incorporates mean versus average deviation. Also, Ferrer et al. (2012) propose a robust optimisation approach which includes demand uncertainty and risk aversion for retail pricing. Mitra and Wang (2005) look at mean-variance, mean-standard-deviation and mean-conditional-value-at-risk approach for deriving a risk-sensitive objective function with revenue management application in traffic and networks. Koenig and Meissner (2010) demonstrate that risk considerations might lead to different decisions when deciding between a quantity-based and price-based revenue model.

In a recent paper, Tang et al. (2012) focus on the risk of the supply side when applying a dynamic pricing strategy. They investigate the newsvendor problem where both yield and demand are random.

Also applying risk considerations to the dynamic capacity model, Huang and Chang (2011) show the effect of using a relaxed optimality condition instead of the optimal one. They investigate model behaviour in numerical simulations and discussed results, given as mean and standard deviation and in a ranking based on a Sharpe ratio. A related approach is presented by Koenig and Meissner (2015), who provide a detailed study of several risk-averse policies for the dynamic capacity model by applying risk measures.

Regarding the use of $V@R$, Lancaster (2003) provides some strong arguments. He demonstrates that risk-neutral revenue management models are vulnerable to the inaccuracy of demand forecasts. Inspired by the $V@R$ metric, he recommends the relative revenue per available seat mile at risk metric. His metric measures the expected maximum of underperformance over a time period for a given confidence level.