Computers & Operations Research 59 (2015) 104–118

Contents lists available at ScienceDirect

Computers & Operations Research

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

Risk management policies for dynamic capacity control

Matthias Koenig [a,1], Joern Meissner [b,*]

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

A R T I C L E I N F O

Available online 6 January 2015

Keywords:
Dynamic decisions
Capacity control
Revenue management
Risk

A B S T R A C T

Consider a dynamic decision making model under risk with a fixed planning horizon, namely the dynamic capacity control model. The model describes a firm, operating in a monopolistic setting and selling a range of products consuming a single resource. Demand for each product is time-dependent and modeled by a random variable. The firm controls the revenue stream by allowing or denying customer requests for product classes. We investigate risk-sensitive policies in this setting, for which risk concerns are important for many non-repetitive events and short-time considerations.

Numerically analysing several risk-averse capacity control policies in terms of standard deviation and conditional-value-at-risk, our results show that only a slight modification of the risk-neutral solution is needed to apply a risk-averse policy. In particular, risk-averse policies which decision rules are functions depending only on the marginal values of the risk-neutral policy perform well. From a practical perspective, the advantage is that a decision maker does not need to compute any risk-averse dynamic program. Risk sensitivity can be easily achieved by implementing risk-averse functional decision rules based on a risk-neutral solution.

1. Introduction

Consider a dynamic decision model under risk for capacity control with a given planning time horizon. The decision maker acts on previous gained information up to a distinct time period and estimations for future time periods. This kind of dynamic decision making under risk is often modelled by dynamic programming formulations. Despite some known limitations of expected utility theory, as discussed by Schoemaker [1], the expected utility approach is often used with dynamic programming for risk considerations. To this end, dynamic programming uses a utility function as an objective function, and time preferences can be included by a discount factor. The books by Chavas [2] and Bertsekas [3] include a description of this approach from a general perspective.

The considered capacity control model is typical, for example, in the area of revenue management, whose use is common in industries such as airlines, hotels or rental cars, in which a firm operates in a monopolistic setting offering multiple products consuming a single resource. The firm owns a fixed capacity of the resource which has to be sold over a finite horizon. The objective of the firm is to find a policy in order to optimise total revenue by allocating capacity to different classes of demand. Usually, a risk-neutral optimisation objective is sufficient for revenue management problems due to the long-term average effect in situations with repeating decision-making processes.

There are, however, many situations when the number of reiterations is too small (e.g., Levin et al. [4] mention an event promoter) or when constraints on working capital or revenue streams force the use of a dynamic decision model with consideration of risk. Weatherford [20] observes that analysts were uncomfortable with risk-neutral objectives and changed waiting lists recommended by their revenue management systems. This means, they applied manually their own risk-averse policy. In practice, short-time objectives of management are a motivation for risk aversion as pointed out by Feng and Xiao [22]. The authors emphasise the obvious effect that uncertainty in demand, forecast and capacity may lead to a significant difference between the realised revenue and expected revenue. The practitioners' demand for risk aversion has motivated the research of risk-averse policies and, thus, this paper, too.

Furthermore, recent approaches by Barz and Waldmann [6] and Huang and Chang [7] propose risk-averse policies for the dynamic capacity control. This model is introduced a standard revenue management model by Lee and Hersh [5] and is originally stated as a dynamic programming formulation of a risk-neutral policy. Barz and Waldmann [6] analyse the dynamic capacity control model under constant absolute risk aversion using an exponential utility as the objective function in the dynamic programming recursion. Huang and Chang [7] present a policy which includes a discount factor not in the objective but in the decision function. This discount factor

* Corresponding author.
E-mail addresses: matthias.koenig@fh-bielefeld.de (M. Koenig), joe@meiss.com (J. Meissner).

[1] Now at: University of Applied Sciences Bielefeld, Germany.

http://dx.doi.org/10.1016/j.cor.2014.12.004
0305-0548/

actually determines a risk premium for certainty of earning revenue now, instead of under uncertainty later. This kind of risk premium is more easily communicated to practitioners than the exponential utility function, where the computation of the risk premium requires certain knowledge about the distribution of the demand function. Huang and Chang [7] also propose a policy considering the selling history and conduct an extensive analysis for risk aversion, comparing standard deviations and Sharpe ratios of risk-neutral and risk-averse policies.

Our objective is the evaluation of a set of control policies under risk considerations. To this end, we perform an analysis of the policies by numerical experiments and look at risk measures in terms of volatility by the standard deviation and in terms of downside risk by the conditional-value-of-risk (CVaR). We extend the analysis of Huang and Chang [7] and propose improved policies which are also easily implemented in practice. Furthermore, we introduce a new straightforward policy which provides acceptable results for moderate levels of risk aversion.

The contribution of this paper is an improvement for applying risk-averse policies. Our presented approach offers advantages for a decision maker in terms of computational and memory requirements. The advantages include less requirements on computing a risk-averse solution and an easy and understandable way of implementing such a solution for practitioners. This is achieved as only one dynamic programming solution is needed for the application of policies of various levels of risk sensitivity. As decision makers often have to determine their level of risk aversion by trying out different levels in simulation, where each level requires a dynamic programming solution, they benefit from our method which requires only the risk-neutral solution for each level. Also, the risk-neutral solution can serve as a basis for applying different levels of risk sensitivity to certain instances in the same setting, i.e., when the risk level is changed from instance to instance. Additionally, we propose a policy which allows to switch risk aversion on or off depending on the current state of the selling rate. This proposed approach could be used with a de(activation) of risk aversion dependent on other possible variables, too.

In particular, we demonstrate that no extra dynamic programming recursions are required for implementing decision rules for risk-sensitive policies. The risk-averse decision can be applied directly using the results of the risk-neutral case. In revenue management terms, it is sufficient to use decision rules directly with the marginal capacity values of the dynamic programming solution of the risk-neutral case.

The remainder of this paper is as follows. Section 2 gives a summary of related work about risk considerations in revenue management context. We describe the dynamic capacity control model, risk-neutral and risk-averse policies associated with the model in Section 3. In Section 4, we present the settings of the numerical experiments and the obtained risk measures evaluating the policies. Finally, we summarise and conclude this paper in Section 5.

2. Related work

A general but comprehensive coverage of revenue management is provided by Talluri and van Ryzin [8] for risk-neutral decision makers. Chiang et al. [9] give an extensive literature overview of the field.

The first revenue management model incorporating risk, the model of Feng and Xiao [10], considers a single-resource problem with two given prices and allows only one price change. They define risk by sales variance as a result of price changes. Their objective function combines expected revenue and a weighted penalty function for sales variance. The weight determines the level of risk aversion. Although their model is limited, the derived result is quite intuitive: risk-averse firms switch to a lower price sooner than risk-neutral ones. This coincides with the risk-averse policies described in Section 3, where firms prefer to accept revenue sooner rather than later.

Lancaster [11] looks at the risk issues in airline revenue management from a practical perspective. He illustrates the vulnerability of revenue management systems by analysing the volatility of historical data of revenue per available seat mile. He runs several simulations which highlight the importance of risk considerations under thin profit margins and high uncertainty. Therefore, he recommends a relative revenue-per-available-seat-mile-at-risk metric which integrates risk measurement with the value-at-risk metric. This metric is the expected maximum of underperformance over a time horizon at a chosen confidence level. To compare different revenue management strategies, he proposes the use of the Sharpe ratio instead of direct dual objective optimisation. This is computationally impractical as revenue distributions are acquired by history or simulations. The arguments [11] for using simulations also hold for our approach of comparing risk measures of different policies for dynamic capacity control.

Risk sensitivity is incorporated by Levin et al. [4] 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 a value-at-risk formulation, and their approach corresponds with maximising expected return subject to a small disaster probability. Risk aversion is introduced in the objective function as a penalty term reflecting the probability that total revenues fall below a certain level. Thus, the underlying utility function at every point in time is piece-wise linear and discontinuous at the point of the desired revenue level.

Discussing risk modelling for traffic and revenue management in networks, Mitra and Wang [12] analyse mean-variance, mean-standard-deviation and mean-conditional-value-at-risk for formulation of the objective function, finally selecting standard deviation as the risk index. The impact of several levels of risk aversion is demonstrated by an efficient frontier for a truncated Gaussian demand distribution.

Koenig and Meissner [13] compare expected revenue and risk in terms of standard deviation and conditional-value-at-risk of pricing policies. A list pricing policy, following capacity control, and a dynamic pricing policy, steadily adjusting prices, are analysed under consideration of the cost of price changes. They show by numerical experiments under which circumstances a policy might be more advantageous over the other.

Robust optimisation [14, cf.] as a means for maximising over a set of worst case outcomes under guaranteed feasibility has been used by various authors in a revenue management context. The worst outcomes are all the smallest revenues under feasible worst-case demand realisations. The works of Perakis and Roels [15], Thiele [16] and Lim and Shanthikumar [17] are exemplary for addressing the problem of uncertainty in the demand function by robust optimisation. Lim and Shanthikumar [17] show that the robust pricing problem is equivalent to a single-product revenue management problem with an exponential utility function without model uncertainty.

Mulvey et al. [18] propose a different approach and consider robustness of solutions in a set of scenarios. They introduce a penalty function to the objective function to achieve a tradeoff between optimality and feasibility. Following this approach, Lai and Ng [19] set up a model for hotel revenue and formulate a tradeoff between expected revenue and mean absolute deviation.

Using expected utility theory in the revenue management context is endorsed by Weatherford [20]. He discusses the assumption of risk-neutrality for a standard revenue management algorithm and concludes that optimising expected utility instead of expected revenue is a suitable risk-averse strategy. In particular, he