A Survey on Risk-averse and Robust Revenue Management

13 April 2017

Jochen Gönsch

Jochen Gönsch

Mercator School of Management,
University of Duisburg-Essen, Lotharstraße 65, 47057 Duisburg, Germany

Phone: +49 203 379 4369
Fax: +49 203 379 1760
E-mail: jochen.goensch@uni-due.de

Abstract

Many industries use revenue management to balance uncertain, stochastic demand and inflexible capacity. Popular examples include airlines, hotels, car rentals, retailing, and manufacturing. The classical revenue management approaches considered in theory and practice are based on two assumptions. First, demand – as the only uncertain variable – follows a known distribution and, second, risk-neutrality justifies the maximization of expected revenue.

Recently, two related streams of literature emerged that do not need these assumptions. Research on risk-averse revenue management acknowledges that, in practice, many decision makers are risk-averse. Research on robust revenue management focuses worst-case scenarios without a known demand distribution, which is especially relevant for new and extremely unstable businesses.

This paper motivates the consideration of risk-averse and robust revenue management. We briefly introduce revenue managements' two main methods – capacity control and dynamic pricing – in the classical, risk-neutral setting. Then, we provide an exhaustive review of the literature on risk-averse and robust capacity control and dynamic pricing. In doing so, the relevant decision criteria are briefly introduced. Finally, possible avenues for future research are outlined.

Keywords: revenue management, capacity control, dynamic pricing, risk-aversion, robustness

1 Introduction

Modern *revenue management* emerged after the deregulation of the US aviation industry in the 1970s. It instantly became a must-have in the airline industry (Cross (1997)). Today, revenue management entails a series of quantitative methods to optimize revenue by controlling sales to heterogeneous customers stochastically arriving over time. The objective is to efficiently use inflexible capacities that are available for only a limited time. Even with revenue management, the revenues will be uncertain, in the classical setting due to the unknown demand.

1.1 Risk-neutral Revenue Management

The classical approach is to use random variables with a known distribution and to maximize the expected revenue, or the profit margin, in industries with considerable variable costs. The high volumes of similar sales processes justify the corresponding *risk-neutrality*. Airlines record hundreds—the major ones even thousands—of departures per day. Given the high level of repetition, a single realization has a negligible impact and the law of large numbers dictates that the average revenue will converge to the expected value. Furthermore, the maximization of the expected value leads to mathematically simpler and more manageable models.

1.2 Risk-averse Revenue Management

Feng and Xiao (1999), followed by Lancaster (2003), first question the assumption of risk-neutrality, but stick to the assumption of known probabilities, leading to *risk-averse revenue management*. Barz (2007) argues that, if the level of repetition is too low, the assumption of perfect insurance and (where applicable) capital markets is necessary to convert the uncertain revenue stream into a certain one with the same present value. Levin et al. (2008) describe the now often-used example of an event promoter who only organizes a few, but very large, events per year. Each event entails a substantial capital investment and could be a matter of economic survival. The promoter would therefore rather have reasonably low revenue expectations and put more emphasis on ensuring that a certain minimum revenue is attained. Optimizing expected values can also be inappropriate when a constant revenue stream should be secured in order to, for example, service financial obligations, or meet shareholders' expectations (Lancaster (2003)). A few negative reports, particularly on the stock market, can often quickly overshadow slightly better average results. Furthermore, decision makers' wishes and preferences are often mentioned. For instance, an advisor for smaller airlines consulted Weatherford (2004) about risk-averse capacity controls. Barz (2007) reports that a consultant's clients were uncomfortable with the results of the risk-neutral system. These clients therefore often change the underlying forecast manually to obtain less risky decisions. Singh (2011) observes that the individual risk-aversion of employees has a significant influence on manual overwriting of decisions from the revenue management system.

1.3 Robust Revenue Management

Robust revenue management assumes that exact distributional information is not available and uses uncertainty sets to describe ranges, for example for demand realizations or distribution parameters. Some worst-case criterion is optimized. While probabilities are usually available in stable business environments with large amounts of historical data, new or unstable businesses often rely on expert judgements (Perakis and Roels (2010)), but also numerical methods are available to construct the uncertainty set from historical data (e.g. Sierag and van der Mei (2016)), think of lower and upper bounds on demand. While robust revenue management is formally different from risk-averse revenue management, many researchers also have more stable revenues in mind and, for example, numerically analyze the expected revenue vs. variance trade-off of robust approaches.

1.4 Sources of stochasticity and uncertainty

In reality, all parameters – demand, capacity, and revenues – can be subject to stochasticity and uncertainty. Environmental factors such as unexpected events, changes in the competitive landscape, or the unexpected impact of alliance partners' flawed revenue management systems might influence the parameters. Capacity might change due to unexpected equipment malfunction or delays. Surprisingly, even the products' prices are uncertain. Especially traditional network airlines have developed very complex fare structures, leading to a vast number of fares on the same route. However, forecasting demand for each fare would be impossible and, even more important, most revenue management systems are still restricted to 24 booking classes (the letters of the alphabet). Thus, fares are grouped and the 'products' used in revenue management are in fact averages over several individual fares. Accordingly, their prices are uncertain as they depend on the actual mix of fares.

In classical, risk-neutral revenue management, a few authors consider variation other than demand variation. For example, de Boer (2004) as well as Wang and Regan (2006) anticipate exogenous aircraft swaps, that is, stochastic capacity in parts of their work. However, their ultimate goal are endogenous swaps based on the development of demand (see also Section 7.6.2).

To the best of our knowledge, stochastics and uncertainty sets always refer to demand in the literature on risk-averse and robust revenue management. Only Lai and Ng (2005) consider uncertain prices (see Section 3.2). Accordingly, this survey is restricted to demand variation as the only source of stochasticity and uncertainty. However, considering potential additional sources might be a possible avenue for future research (see Section 7.6.3).

1.5 Contribution

Although much called for in practice, no survey or textbook has covered this active area of research that started about two decades ago. This paper attempts to close this gap by providing an exhaustive literature review along with an introduction into revenue management. It aims at readers with a general interest in revenue management as well as risk-aversion and robustness, but does not presume any prior knowledge. Finally, we compare the state of the art with recent trends in related areas and give an extensive outline of possible avenues for future research.

1.6 Outline

The remainder of this paper is structured as follows: In Section 2, we introduce classical, expectation-maximizing revenue management and briefly explain its two major areas, capacity control and dynamic pricing with the corresponding basic model formulations. Sections 3 and 4 review the literature on risk-averse and robust capacity control, respectively. Sections 5 and 6 summarize the corresponding literature on dynamic pricing. In Section 7, we outline potential avenues for future research.

2 Classical Revenue Management

Historically, as well as methodologically, one can distinguish between two basic approaches. *Capacity control* deals with products offered at prescribed prices. In passenger aviation, these are the different booking classes comprising an identical core service (a seat on a flight from A to B in economy class) with different conditions (options to change or cancel the booking, price, etc.). In Germany, another well-known example are the national railway company's budget prices (€19, €29, etc.; e.g. Nahler (2013)). In contrast, *dynamic pricing* influences demand by adapting prices over time. Examples include low-cost carriers like Ryanair (e.g. Malighetti et al. (2009)) or the end-of-season sales in fashion retailing. Mostly, non-negotiable posted prices are used.

2.1 Capacity Control

In the following, we introduce the stochastic, dynamic standard model of capacity control. We limit ourselves to the marketing of a single resource (the single-leg setting in airline revenue management) and independent demand, which presumes that each product corresponds to exactly one disjoint customer segment. For a detailed account of the fundamentals and the marketing of multiple resources (e.g., connecting flights, a network), as well as how customer choice behavior is taken into account, we refer to the literature. The textbook by Talluri and van Ryzin (2004, Chapters 1-3), as well as the recent description in Gallego et al. (2015), provide a good overview and further references. Besides airlines, research on the applications of capacity control includes, among others, car rental companies (e.g. Geraghty and Johnson (1997), Haensel et al.

(2012)), tour operators (e.g. Klein (2000), Xylander (2003)), passenger railways (e.g. You (2008), Kellermann and Cleophas (2015), Hetrakul and Cirillo (2015) and manufacturing (e.g. Spengler et al. (2007), Defregger and Kuhn (2007), Hintsches et al. (2010), Volling et al. (2012), Guhlich et al. (2015) and Seitz et al. (2016)).

The company markets a resource with an initial capacity of C units. Products $i \in \mathcal{I} = \{1, \dots, n\}$ are offered at a price of r_i, and each unit sold requires one capacity unit from the resource. The sales horizon comprises T micro periods (numbered in descending sequence), the length of which is chosen in such a way that at most one customer request arrives in each micro period $t = T, \dots, 1$. The probability of a request for product i in period t is denoted by $p_i(t)$. Thus, the probability that there will be no product request in period t is $p_0(t) = 1 - \sum_{i \in I} p_i(t)$. The maximization of the expected revenue is formulated as a stochastic dynamic optimization problem by using the *Bellman equation*:

$$V(c,t) = \sum_{i \in \mathcal{I}} p_i(t) \cdot \max\{V(c,t-1), r_i + V(c-1,t-1)\} + p_0(t) \cdot V(c,t-1) \quad (1)$$

for all $0 \leq c \leq C$ and $t = T, \dots, 1$. The *value function* $V(c,t)$ denotes the optimal expected revenue from period t onwards with remaining capacity c. The following *boundary conditions* are required: $V(c,t) = -\infty$ for $c < 0$ and $V(c,0) = 0$ for $c \geq 0$.

Whenever a request arrives, a decision about its acceptance has to be made. This is analogous to a decision in respect of the offered products. More specifically, a request is accepted if and only if $r_i + V(c-1,t-1) \geq V(c,t-1)$ or, equivalently, $r_i \geq V(c,t-1) - V(c-1,t-1)$. The term on the right-hand side denotes the *opportunity cost* of accepting a request in period t with remaining capacity c. This means that the request may be rejected in the expectance that the required capacity unit can later be sold at a higher price (as part of a different product). We thus obtain a decision rule, or policy, π, which determines the products to offer in each state (c,t) of the system.

2.2 Dynamic Pricing

The textbooks of Talluri and van Ryzin (2004, Chapter 5), as well as Philips (2005, Chapter 10), present the fundamentals of dynamic pricing in detail. Numerous review papers classify the existing publications. Bitran and Caldentey (2003) as well as Chiang et al. (2007) provide classical literature reviews on dynamic pricing in general. More recent review papers focus on current, intensively researched aspects. Gönsch et al. (2013) provide an overview on dynamic pricing with strategic customers who anticipate that the supplier may lower the price. Den Boer (2015) summarizes the current research status of dynamic pricing with learning. In this area, the firm constantly updates its demand forecast based on past sales. Application areas of dynamic pricing include, for example, retailing (e.g. Zhao and Zheng (2000), Heching et al. (2002)), low cost airlines (e.g. Marcus and Anderson (2008)), hotels (e.g. Schütze (2008)), and make-to-order manufacturing (e.g. Hall et al. (2009)).

The fundamental model in dynamic pricing has many similarities with capacity control. The company in question markets a given stock C of a single product with a sales horizon of length T. The offer price r_t can be adjusted at the start of each micro period $t = T, \dots, 1$ based on the current residual capacity $0 \leq c \leq C$. A sale will occur with probability $p(r_t)$. The maximization of the expected revenue is formulated as a stochastic dynamic optimization problem using the Bellman equation:

$$V(c,t) = \max_{r_t} \left(p(r_t) \cdot \left(r_t + V(c-1, t-1) \right) + \left(1 - p(r_t) \right) \cdot V(c, t-1) \right) \tag{2}$$

for all $0 \leq c \leq C$ and $t = T, \dots, 1$. The value function $V(c,t)$ again denotes the optimal expected revenue from period t onwards with remaining capacity c and the following boundary conditions are required: $V(c,t) = -\infty$ if $c < 0$ and $V(c,0) = 0$ if $c \geq 0$. The price r_t can be chosen from a discrete or continuous set of possible prices. For technical reasons, this set usually must include a sufficiently high price for which there will be no requests. In dynamic pricing, the policy π determines the price r_t for each state (c,t).

3 Risk-averse Capacity Control

There are two frequently used approaches to deal with risk-aversion: expected utility theory and risk-measures. Table 1 (left part) provides an overview of the research published. In the following subsections, we discuss the literature using these approaches in detail. In each subsection, we first introduce relevant criteria and then present the capacity control literature using them.

Risk-averse Capacity Control			Robust Capacity Control		
Expected Utility	Mean-Risk	Risk Measures	Maximin	Regret	Competitive Ratio
Weatherford (2004): additive	Lai/Ng (2005): mean abs. deviation	Gönsch/Hassler (2014): CVaR	Birbil et al. (2009)	Lan et al. (2008): absolute	Lan et al. (2008)
Barz (2007): additive, general atemporal, exponential		Koenig/Meissner (2015a): VaR	Perakis/Roels (2010)	Lardeaux et al. (2010): absolute	Ball/Queyranne (2009)
Barz/Waldmann (2007): atemporal exponential		Koch et al. (2016): CVaR*	Rusmevichientong/ Topaloglu (2012)	Gao et al. (2010): relative	Lan et al. (2015)
Feng/Xiao (2008): atemporal exponential		Koenig/Meissner (2016): TPR	Sierag/van der Mei (2016)	Perakis/Roels (2010): absolute	Gao et al. (2016)
Zhuang/Li (2011): general atemporal				Lan et al. (2011): relative	
Koch et al. (2016): exponential*				Lan et al. (2015): absolute	
Huang/Chang (2011): no clear objective; STD and sharp ratio for evaluation					
Koenig/Meissner (2015b): no clear objective; STD and CVaR used for evaluation					(*arbitrary objectives)

Table 1: Overview of research on risk-averse and robust capacity control

3.1 Utility Functions in Capacity Control

Expected utility theory developed by von Neumann and Morgenstern (1944) derives the existence of a non-decreasing *utility function* $u: \mathbb{R} \to \mathbb{R}$ from simple axioms. Let the random revenues R_{π_1} and R_{π_2} result from two different policies $\pi_1 \neq \pi_2$. The decision maker is risk-averse if a deterministic payment equaling the expected value causes a higher utility than the expected utility: $u\big(\mathbb{E}(R_\pi)\big) \geq \mathbb{E}\big(u(R_\pi)\big)$ for all π. One can therefore derive, with help of Jensen's inequality, that a decision maker with a utility function u is risk-averse when u is concave. A *linear utility function* implies *risk-neutrality,* and a *convex utility function* expresses *risk-seeking.*

In multi-period models, the aggregation of utility across the periods must be determined. With *additive time-separable utility functions* the total utility $u(R_\pi) = \sum_{t=T}^{1} u_t(R_\pi^t)$ is the sum of the utility functions $u_T, \dots, u_1$ for the individual periods. *Atemporal utility* *functions* model a decision maker who is indifferent between reward streams of the same sum. Thus, only the sum $\sum_{t=T}^{1} R_\pi^t = R_\pi$ of the period revenues R_π^t is relevant and $u(R_\pi) = u(\sum_{t=T}^{1} R_\pi^t)$. This is used, for example, if the planning horizon is short and allows an arbitrary redistribution of the revenue between the individual periods without changing utility. There are two classes of atemporal utility functions that evaluate different alternatives independent of the decision maker's current wealth and thus allow a decomposition by time periods. Obviously, a *linear (risk-neutral) utility function* is additive time-separable. In addition, the frequently used *exponential utility function* $u^\gamma(R_\pi) = -e^{-\gamma R_\pi}$ with constant absolute risk inclination γ allows the decomposition $u^\gamma(R_\pi) = -\prod_{t=T}^{1} -u^\gamma\ (R_\pi^t)$. Here, a positive value of γ implies risk-aversion.

In practice, the determination of utility functions is unfortunately often difficult due to the concept's theoretical complexity and the abstract representation in utility values (e.g., Gotoh and Takano (2007)), requiring the decision maker to answer a multitude of questions (e.g., Xu (2010)).

The first contribution to risk-aversion in capacity control stems from Weatherford (2004). It is not based on the standard model described in Section 2.1, but leverages the classical EMSR (Expected Marginal Seat Revenue) heuristic, which is still widely used in industry. To incorporate risk-aversion, he simply replaced the product's revenues with their utilities and obtained the EMSRU (EMSR utility) rule. While the original EMSR method of Belobaba (1989) clearly tries to maximize the expected revenue—and yields the optimal solution under certain conditions—, the intuitive EMSRU alternative pursues no clear target criterion. The resulting policies, although more conservative than with EMSR, are clearly more risk-seeking than the employed utility function suggests. This is because the transition from total revenue to a marginal consideration is the basis of EMSR but no longer possible in the case of risk-averse utility functions (see also the discussion in Barz (2007, p. 122-123)).

Barz and Waldmann (2007) were the first to publish work on a risk-averse utility function in the stochastic dynamic standard model of capacity control (Section 2.1). They employe the *atemporal exponential utility function* $u^\gamma(R_\pi) = -e^{-\gamma R_\pi}$. By using the abovementioned reformulation, the utility function is integrated into the Bellman equation. The authors show that all the attractive structural properties of the optimal policy that are well-known in the risk-neutral case carry over. It is, therefore, possible to describe an optimal policy by means of protection levels indicating the number of capacity units which a low-value product cannot use. In her PhD thesis, Barz (2007) extends some results to choice-based capacity control, where selling probabilities depend on the set of products offered to a customer and also considers the popular static setting, where low value demand arrives strictly before high value demand. Based on her numerical results, Barz (2007) also conjectures that the optimal protection levels are monotone in the degree of risk-aversion γ. Apparently independently and in a paper submitted already in 2004, Feng and Xiao (2008) pursued a comparable objective and also employ atemporal exponential utility functions. They motivate these particular function type by showing that it essentially consists of the expected revenue and the higher-order moments of revenue, which represent risk. Rather unusual for capacity control, the authors work in continuous time. Their findings are consistent with Barz' (2007) monotonicity conjecture.

Regarding *general concave (risk-averse) atemporal utility functions*, only an evaluation of the cumulative revenue R_π is possible in the last period. To do so, the state space of the Bellman equation must be expanded to include the already obtained revenue z. Thus, the value function $V(c, z, t)$ equals the optimal expected utility for the complete horizon, given a remaining capacity of c and an already obtained revenue of z at t. Unfortunately, the expansion of the state space generally increases the solution complexity and tractability considerably. Things are further complicated by the fact that some intuitively expected monotonicities (e.g. in the remaining capacity) do not hold (Barz (2007)).

Zhuang and Li (2011) focus a static model with only two products (booking classes) and general concave atemporal utility fuctions. Their development is closely related to Eeckhoudt et al (1995), who consider a risk-averse newsboy, and Barz (2007). However, they extend the results in Eeckhoudt et al. (1995) to a static two-class revenue management problem. Barz' monotonicity conjecture is partly resolved by showing that the booking limit is increasing in the degree of risk-aversion in the two-class problem. They then identify a general sufficient condition assuring that a risk-averse manager's optimal booking limit is monotone in the level of demand uncertainty measured by stochastic dominance. With normally distributed demands, they obtain insights on the impact of demand uncertainty and the risk attitude through a mean-variance analysis.

Finally, Barz (2007) also considers *concave additive time-separable utility functions*. Tractability is no issue as basically only r_i has to be replaced by $u_t(r_i)$ in (1). However, the use of such a function is hard to motivate, since the decision maker usually is only concerned about overall revenue and not about its distribution to the (arbitrary small) time periods. Again, some monotonicities carry over, some do not (e.g. protection levels are not monotone in time).

3.2 Mean-risk in Capacity Control

Mean-risk approaches originate from the finance sector. They weigh two criteria against each other: the mean (the expected value) and the risk (a scalar measure for variability). One therefore obtains many efficient expected value/risk combinations, which can be determined by maximizing the expected value for a given risk level (or vice versa). This approach has many advantages, because both criteria can be weighted with a specified parameter and the trade-off can be analyzed. In this context, the variance that Markowitz (1959) used to describe risk in portfolio optimization became very well known. However, like the mean absolute deviation (the average of the absolute deviations from a central point, usually the mean), it can contradict stochastic dominance and, thus lead to unintuitive results (Ogryczak and Ruszczynski (2002)).

Lai and Ng (2005) focus on the hotel industry. The paper is unique as the authors also directly consider uncertain prices. Their stochastic model is based on demand scenarios with given probabilities that differ in the products' prices and demand. The main idea seems to be to penalize the mean absolute deviation of the revenue from its mean with a parameter in the objective function. The only decision variable is the number of accepted bookings (as opposed to the usual maximum number, called partitioned booking limits) and does not depend on the scenario. Thus, revenue deviations are only caused by price variation (and not demand). Moreover, feasibility is jeopardized when demand is less than the number of accepted bookings. These "constraint violations" are penalized with additional parameters, but they symmetrically penalize when not all demand is accepted. It remains unclear why revenue is not calculated from the number of bookings actually accepted in each scenario (minimum of booking limit and demand). The paper is not related to prior or later work, but often cited by from researchers in its specific field of application.

3.3 Risk Measures in Capacity Control

During the past years, numerous risk measures found broad acceptance, particularly in the finance sector, where they are used for ascertaining a bank's risk position and thus target losses. Here, we introduce these risk measures directly in the context of the maximization objectives considered. An important advantage of risk measures—especially with regard to utility functions—is that they are mostly measured in the same (monetary) unit as the underlying random variable. Owing to their comparatively simple definition, decision makers in the industry without a distinct theoretical background can also understand risk measures easily. It is important to note that the random variable the risk measure is based on is usually cumulative revenue over the entire booking horizon.

In the mid-1990s, J.P. Morgan introduced the *value-at-risk (VaR)* indicator in the product RiskMetrics to estimate the aggregated risk of all positions of a bank. Today, VaR is widely used in practice. The VaR at level α measures the revenue that will not be exceeded with a probability of α. Formally, it corresponds to the quantile function (inverse of the distribution function): $\mathrm{VaR}_\alpha(R_\pi) = F_{R_\pi}^{-1}(\alpha) = \inf\{x: F_{R_\pi}(x) \geq \alpha\}$. The VaR is positively homogeneous, monotonic, and translation invariant, but generally not subadditive. It thus contradicts the idea of diversification. From an intuitive perspective, it is often criticized because the distribution of the revenue below the quantile is not taken into consideration. *First-order stochastic dominance* is equivalent to a *VaR at least as high for each level α* (Ogryczak and Ruszczynski (2002)).

The *conditional value-at-risk (CVaR)* stems from a pioneering article on risk measures by Artzner et al. (1999). They define desirable characteristics (positive homogeneity, monotonicity, translation invariance, and subadditivity) of risk measures, which led to the *coherent risk measures* concept. For continuous distributions, the CVaR for level α corresponds to the conditional expected value: $\mathrm{CVaR}_\alpha(R_\pi) = \mathbb{E}\big(R_\pi \,\big|\, R_\pi \leq \mathrm{VaR}_\alpha\big)$. More complex, but manageable representations exist for discrete distributions. Different from the VaR, the CVaR takes the distribution below the critical quantile (of VaR) into account. A clear interpretation of CVaR_α is: "What average revenues will be achieved if revenue is below the VaR_α?" As CVaR is subadditive, it rewards diversification. Owing to its intuitive and theoretical advantages, CVaR is now increasingly used in theory and practice. *Second-order stochastic dominance* is equivalent to a *CVaR at least as high for each level* α (Dentcheva and Ruszczynski (2006), Noyan and Rudlof (2013)).

Recent publications also propose the *target-percentile risk* (*TPR*, see Boda and Filar (2006)). Here, the goal is to exceed a given target value z (satisfaction target) and the probability not to exceed this target is measured: $TPR(z) = P(R_\pi \leq z)$. Hence, this probability is minimized. *First-order stochastic dominance* is by definition equivalent to a *TPR at least as high for each target value* z.

The first risk measure optimized in capacity control was the *TPR* around 2010, although Koenig and Meissner (2016) published their paper much later. Analogous to the case of atemporal utility functions, the state space must be expanded here. The value function $V(c, z, t)$ then depends on the remaining capacity c, the revenue z still to be achieved, and the time t. It denotes the probability of not exceeding a cumulative revenue of z in the remaining periods $\tau = t, \ldots, 1$. Therefore, in this formulation, z stands for the revenue that must be exceeded in periods $\tau = t, \ldots, 1$. Accordingly, it is reduced with the acceptance of a request.

Koenig and Meissner (2015a) were the first to maximize the *VaR*. They build on their paper on the TPR and employ the following relation: $\mathrm{VaR}_\alpha = \sup\{z \in \mathbb{R}: P(R \leq z) < \alpha\} = \sup\{z \in \mathbb{R}: \mathrm{TPR}(z) < \alpha\}$. By means of a binary search, they determine the highest target value z for which the $\mathrm{TPR}(z)$ is still below the level α.

An integration of the *CVaR* into the Bellman equation, and therefore an analytical optimization, appears impossible due to the discrete decision space in capacity control (Gönsch and Hassler (2014)). However, based on the approach of Pflug and Pichler (2016) for convex decision spaces, the authors developed an efficient heuristic. With their results, one can illustrate the fundamental trade-off between a high expected revenue and low fluctuations (Figure 1). While the expected value (EV) is maximized for $\alpha = 1$, the average revenue decreases with increasing risk-aversion (smaller α), until it eventually equals the revenue at the acceptance of each request (FCFS). It is interesting to see that the standard deviation (and, hence, the variance) does not decrease continuously.

3.4 Other Approaches in Capacity Control

While the papers mentioned thus far all focus on the optimization of a well-defined target criterion, some papers cannot be clearly categorized. Neither Huang and Chang (2011), nor Koenig and Meissner (2015b), follow a clear target criterion. Huang and Chang (2011) developed intuitive heuristics based on the stochastic standard model. More specifically, they extend the classical policy, where a request is accepted only when the resulting revenue exceeds the opportunity costs ($r_i \geq V(c, t-1) - V(c-1, t-1)$) with a risk factor $\delta \in [0,1]$, and obtain $r_i \geq \delta\big(V(c, t-1) - V(c-1, t-1)\big)$. A more conservative policy results for $\delta < 1$, which tends to accept more requests with lower revenue. To adjust the risk factor over time, Huang and Chang recommend various intuitive formulas that are manually parameterized. Koenig and Meissner (2015b) develop this approach further and also incorporate the risk factor in the Bellman equation. Koch et al. (2016) embed the approaches in a framework to specifically optimize a target criterion. To

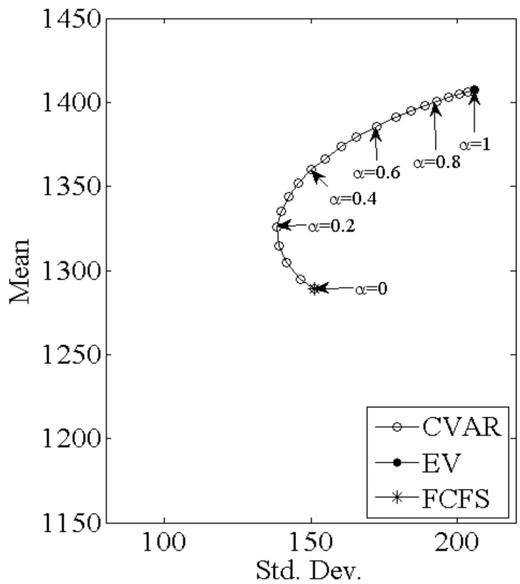

Figure 1: Mean and standard deviation (optimization of the CVaR)

do so, they determine the parameters of the risk factor formula with simulation-based optimization, following the risk-neutral approach of Klein (2007). While most papers consider specific simplified problems, this approach is applicable to almost any problem and objective. They also consider customer choice behavior, which only Barz (2007) considers otherwise.

4 Robust Capacity Control

Robust capacity control is employed when no or only limited distributional information is available, for example, when a new business or product is launched or demand is difficult to predict. As the decision criteria described in the previous section are not applicable, various criteria have been developed specifically for robust decision making. They all have in common that some worst-case scenario is optimized. However, most authors also have in mind to employ these approaches to reduce revenue variability. Thus, often an expected revenue vs. variance trade-off is examined in numerical experiments. Please note that we do not consider approaches here that are model-free or distribution-free but aim at maximizing expected revenue. For example, literature on simulation-based and stochastic optimization (e.g. Gosavi et al. (2007), Topaloglu (2008), van Ryzin and Vulcano (2008), and Kunnumkal and Topaloglu (2009)) does not assume any specific demand or cancellation distribution but usually only requires that samples from the demand distribution are available.

Please note that the term *robust optimization* is often associated with the question of feasibility in the sense that the chosen decision should be feasible for every environment state. Feasibility in this sense is, however, irrelevant in revenue management, because policies already contain rules that only allow a sale if there is sufficient capacity, independent of the optimization results. Nonetheless, we use the term ‘robust’ in line with the vast majority of the literature. Only two papers emphasize the lack of distributional information by using the terms ‘limited demand information’ (Lan et al. (2008)) and ‘distribution-free’ (Gao et al. (2016)) in their titles.

In the following, we briefly introduce the criteria widely-used in robust revenue management. Then, two subsections discuss the relevant research in detail. Table 1 (right part) provides an overview. For literature on decision theory, robust decision making, and use of regret in various operations management problems we refer the reader to Perakis and Roels (2008), Kouvelis and Yu (1997), and the references therein. Central to all criteria is the set of possible environment states (uncertainty set).

The *maximin criterion* prefers the policy π with the highest worst-case revenue. This surmises an *extreme risk-averse* attitude, because only the worst scenario is considered.

The *regret-based criteria* consider the difference (regret) of the revenue obtained in a scenario with a policy and the revenue obtained with a policy that is optimal for the scenario. This optimal policy is determined with complete information (perfect hindsight or clairvoyance) and is often denoted as *ex-post* in revenue management. The *absolute minimax regret criterion* (also Savage-Niehans criterion) focuses on the absolute difference between revenues. In contrast, the *relative minimax regret criterion* considers the ratio of the regret to the optimal (ex-post) revenue.

The *competitive ratio* formally resembles the relative minimax regret criterion. However, it stems from the analysis of online algorithms (see, e.g., Albers (2003)). An online algorithm addresses a sequential decision-making problem. The algorithm (corresponding to a policy) processes an input stream (customer requests for the products) and makes an immediate decision for each input (request), without knowledge of future inputs. Its counterpart is a competitor, the offline algorithm, which has knowledge of the entire input sequence and only decides ex-post (the optimal policy). It is equivalent to assume that the competitor also decides immediately, but he can determine the input continuation.

4.1 No Information regarding the Demand Distribution

Ball and Queyranne (2009) are the first to use techniques from the analysis of online algorithms in single-leg (aka single-resource) capacity control. They derive static, nested booking limits and consider bid-price controls. Their policies come with guarantees regarding the *competitive ratio* without any information about demand. For the static problem, they obtain closed-form optimal solutions.

Lan et al. (2008) develop an integrated approach for the consideration of the *absolute minimax regret criterion* and the *competitive ratio.* For both criteria, the optimal static policies are nested booking limits. Dynamic booking limits are also considered. They generalize the aforementioned paper by considering lower and upper bounds on demand for each product and provide alternative proofs for some of their results. The optimal policies contain those of Ball and Queyranne (2009) as a special case.

Perakis and Roels (2010) are even more general in some respects as they consider network revenue management with a general polyhedral uncertainty set and various static controls: partitioned booking limits, nested booking limits, displacement-adjusted virtual revenues, and bid-prices. The authors describe in detail how expert knowledge on demand can be used in the construction of the set of environment states and consider the *maximin criterion*, as well as the *absolute minimax regret criterion.* They find that the minimax regret approach performs very well on average and—despite its worst-case focus—outperforms standard controls when demand is correlated or not unconstrained. The maximin approach is more conservative but provides a revenue guarantee.

Gao et al. (2010) consider competition using the *relative minimax regret criterion* and thereby indirectly also the competitive ratio without any information on demand. They focus on a static model with nested booking limits for two parallel, single-leg flights with only two products each, offered by two competing airlines and show the existence of a Nash equilibrium, as well as the value of a coordination.

With Gao et al. (2016), the authors of Lan et al. (2008) extend their paper to time-variant information on demand. More precisely, time is sliced into several time periods and there are now bounds on demand for each product in each time period. These time periods should not be confused with the micro periods of (1). The authors consider *competitive ratio* in settings with and without updating of booking limits during the booking horizon. Comparing the former with the results from Lan et al. (2008) shows the benefit of using additional (temporal) information on demand. According to the numerical results, policies are more flexible and show a superior performance. While optimal booking limits can be efficiently derived from a MIP in the static setting, this is not possible in the dynamic setting because structural properties no longer hold. Thus, high-quality heuristics using closed-form solutions are provided.

In the context of robust capacity control, Lan et al. (2011) are the first to incorporate overbooking decisions. They consider a single-leg, multiple-product setting and focus on *relative regret*. Following Lan et al. (2008), demand is characterized by bounds for each product. Analogously, the share of passengers that request service (show up at departure) is given by upper and lower bounds. The model developed has appealing analytical properties: the worst case scenarios are characterized, and their number is only $n + 1$ with n products. This allows to obtain the optimal solution in closed form. The optimal static policy retains the well-known structure of nested booking limits. Finally, the overbooking level obtained is an alternative to the existing static, heuristic overbooking methods. It can be coupled with existing capacity control approaches without overbooking like EMSR that ignore cancellations.

In a follow-up paper, Lan et al. (2015) consider ‘hybrid’ information on demand, that is, demand is characterized by bounds, but distributional information regarding no-shows is available. They use the *competitive ratio* and the *minimax absolute regret criterion* and show that nested booking limits are optimal among all deterministic policies. Based on structural properties, the worst-case scenarios to consider are characterized, but closed-form solutions for the nonlinear stochastic optimization problem are not obtained. Thus, a model-based heuristic is provided. Again, the developed overbooking approach can be combined with existing capacity control approaches.

Lardeux et al. (2010) stand out as they build on a deterministic approach and are primarily concerned with capturing stochastics. More precisely, they extend the Certainty Equivalent Control (CEC) which decides on requests’ acceptance based on opportunity cost approximated by the difference of two models. While stochastics is considered in the standard RLP (Talluri and van Ryzin (1999)) approach via the averaging of solutions of independent instances for different scenarios regarding total demand, the authors minimize *absolute regret* of common partitioned booking limits over all scenarios. Due to the CEC approach, this calculation is repeated on every request's arrival, leading to a considerable computational burden. From a theoretical point of view, it is interesting to observe that, thus, regret inherently relates only to the remaining booking horizon (see also the discussion on time consistency in Section 7.4).

4.2 Incomplete Information regarding the Demand Distribution

The second approach to capacity control under uncertainty assumes that demand, in principle, follows certain distributions, but they are not fully known. Birbil et al. (2009) consider static and dynamic single-leg models. In the static model, the distribution of the amount of demand for individual products is tainted with a certain vagueness. Thus, the set of potential environment states entails a continuum of discrete distributions for each product. In the dynamic model, the arrival probabilities ($p_i(t)$ in (1)) are not exactly known, but belong to an ellipsoidal uncertainty set. The authors consider the *maximin criterion* and provide the Bellman equation. Numerical experiments show that the approach achieves a reduction in variability (about 1%-20% less standard deviation) at only a slight decline in average revenue (up to 1.5%) in comparison to a standard, non-robust control in static and dynamic settings where the parameters of the distribution are not exactly known.

Rusmevichientong and Topaloglu (2012) also use the *maximin criterion*, but consider customer choice. The parameters of the multinomial logit (MNL) model belong to an uncertainty set. As usual, the size of this set reflects the degree of uncertainty. Alternatively, the size can be interpreted as a means to control the trade-off between increasing the average revenue and protecting against the worst-case scenario. In contrast to static revenue management models, their static model does not consider capacity constraints. Thus, it is a robust, parametric version of the assortment optimization problem and complements, for example, the nonparametric approach to calculate (but not optimize) revenue from an assortment in Farias et al. (2013). The dynamic model considers the standard single-leg setting from revenue management and enhances the stochastic dynamic model (1) with customer choice. The important nesting by fare order property carries over. Numerical experiments show that worst case performance indeed improves while average revenues are slightly smaller.

Sierag and van der Mei (2016) consider the standard, dynamic single-leg setting with customer choice and cancellations. They directly work with the choice probabilities of arbitrary choice models, which lie in an ϕ-divergence uncertainty set and the *maximin criterion* is used. According to the authors, a major unexpected numerical finding is that robust methods perform better for smaller inventories. However, no further investigation is provided and this might also be due to other factors like a change in demand variability (constant number of time periods with different total demands). In addition, also the benchmark heuristic's performance considerably deteriorates in inventory size (from -1% to -11%). Moreover, the results suggest that the approach performs especially well if no information about cancellations is available.

5 Risk-averse Dynamic Pricing

Table 2 provides an overview on the literature on risk-averse and robust dynamic pricing. Analogously to capacity control, risk-aversion in dynamic pricing is mostly incorporated via utility functions and newer papers also consider risk-measures. One criterion that is not considered in the context of capacity control is also used.

Risk-averse Dynamic Pricing				Robust Dynamic Pricing	
Expected Utility	Chance Constraints	Mean-Risk	Risk Measures	Maximin	Regret
Lim/Shanthikumar (2007): exponential	Levin et al. (2008): minimum revenue	Feng/Xiao (1999): variance	Gönsch et al. (2017): CVaR	Lim/Shanthikumar (2007)	Wang/Xiao (2017): absolute
Lim et al. (2008): exponential			Schur et al. (2017): (CVaR)	Lim et al. (2008)	
Li/Zhuang (2009): additive, atemporal exponential					
Schlosser (2015, 2016): exponential					

Table 2: Overview of research on risk-averse and robust dynamic pricing

Compliance with an *aspiration level* (see, e.g., Simon (1959)) can be enforced with constraints (see, e.g., Prékopa (2003) for an overview regarding optimization). *Chance constraints*, which, for example, limit probabilities of adverse events, are a customary tool in optimization. They can be viewed as a relaxed version of stochastic dominance. Many of the criteria described in Section 3 can also be applied to formulate chance constraints.

Feng and Xiao (1999) were the first to focus risk-averse dynamic pricing. They consider the case of two preset prices and investigate the optimal moment for a one-time possible switch between the two. Their *mean-risk approach* penalizes revenue variation by means of a risk factor in the objective function. The continuous-time model is solved in closed form and the optimal policy is analyzed. Among others, a more risk-averse firm switches earlier to the lower price, all else being equal.

Levin et al. (2008) maximize expected revenue in the objective function, subject to a *chance constraint* which ensures that a given minimum revenue is attained with a given probability. As decision makers often do not know the exact values for the minimum revenue and the probability, the authors suggest to construct an efficient frontier in the plane of probability and expected revenue. This frontier can be alternatively obtained by penalizing the probability of falling short of the minimum revenue in the objective function and varying the penalty factor. The authors formulate the model as a continuous-time optimal control problem, obtain optimality conditions, and analyze structural properties of optimal policies. Approximation techniques are provided as the optimal dynamic pricing policy may be hard to compute because of the size of the state space and the large number of steps required for time discretization.

Li and Zhuang (2009) consider *additive time-separable exponential utility functions* and *atemporal exponential utility functions* in discrete time. The authors show that well-known monotonies regarding time and capacity still hold. The optimal price decreases in the risk-aversion in respect of both utility functions all else—including opportunity cost—equal.

Schlosser (2015) endogenizes the decision about the advertising intensity, which, just as the sales price, influences sales. The author succeeds in specifying optimal closed-form policies for *exponential utility functions* in continuous time. This is the first time optimal closed-form solutions are derived for risk-averse dynamic pricing, with the exception of the highly restricted setting of Feng and Xiao (1999).

Schlosser (2016) considers multi-product dynamic pricing in continuous time. The products have independent demands and inventories, but are related through adoption effects. That is, sales of one product can depend on past sales of all products. In an extension, he also considers risk-aversion via an *exponential utility function*. Simulations show that the variance can be significantly reduced, while expected profits are still near optimal. Furthermore, it turns out that the expected evolutions of the risk-averse prices are downshifted versions of their risk-neutral counterparts, but their overall shape remains the same. By contrast, the expected evolutions of optimal risk-sensitive advertising rates are clenched versions of the corresponding risk-neutral path. Hence, he derives the following rule of thumb: Under risk aversion, a suitable absolute mark-down on optimal risk-neutral prices should be combined with a fixed proportion of optimal risk-neutral advertising rates.

With *CVaR*, Gönsch et al. (2017) are the first to optimize a risk measure in dynamic pricing. The sale of a single-capacity unit is analytically solved in discrete time and various structural characteristics are shown. While it is well-known that the optimal price continually decreases in the risk-neutral case, it is optimal for the risk-averse decision maker to start with a low price and to initially keep it constant until, at one point in time, the risk-averse price equals the risk-neutral price. Then, one should follow a risk-neutral policy. This behavior is also often observed in practice, but was to date explained by means of price adjustment costs or similar concepts. Numerical methods are developed for larger start capacities. Computational results show that applying a risk-averse policy, even a static one, often yields a higher CVaR than applying a dynamic, but risk-neutral, policy. The difference is bigger for smaller capacities.

Schur et al. (2017) also consider CVaR, but—in contrast to the aforementioned paper—they do not consider CVaR of total revenue but employ a so-called nested formulation, that simply replaces the expectation in (2) with CVaR. This has some theoretical advantages and disadvantages (see the discussion on time consistency in Section 7.4) and is much more tractable. They show that the risk-averse dynamic pricing problem can be easily transformed to a risk-neutral one. Only the selling probabilities used in (2) have to be slightly adapted. Thus, all properties and monotonicities carry over and existing algorithms and software systems can be reused. While Schlosser (2016) suggested a constant absolute markdown on risk-neutral prices as a rule of thumb in his setting, Schur et al. (2017) analytically show when a fixed proportion is optimal.

6 Robust Dynamic Pricing

Lim and Shanthikumar (2007) consider robust dynamic pricing in continuous time. The authors characterize the demand by means of relative entropy, a distance measure for distributions and model a game against an adversarial nature (*maximin criterion*). Closed-form solutions for an exponential nominal demand are derived and conditions under which this robust dynamic pricing problem is equivalent to the maximization of an *exponential utility function* with a given demand distribution are shown.

Whereas all papers on dynamic pricing mentioned so far consider a single product, Lim et al. (2008) extend the aforementioned paper to multiple products which can have different levels of ambiguity. If demand functions and resources consumed are independent, the problem can be decomposed into risk-averse single-product problems where the certainty equivalent of an exponential utility function is maximized. The magnitude of the risk-aversion parameter depends on the level of ambiguity concerning the associated product's demand. If resources are jointly used, an equivalent risk-averse dynamic pricing problem maximizes the sum of the certainty equivalents of the product's revenue streams. It requires a revenue sharing function that stipulates the transfer of revenue in the instant of a sale to other product's streams that suffer from the resulting decrease in available resources. Finally, a decomposition to single-product problems is provided, albeit still requiring independent demands.

Wang and Xiao (2017) consider *absolute regret* in continuous time with a discrete set of predefined price points. They assume that the demand rate for each price point follows an unknown distribution with finite bounds and only upper and lower bounds for each rate are known. Analogous to robust capacity control, numerical experiments show a good variation vs. average revenue trade-off.

7 Future Research

In this section, we develop and structure avenues for future research. Therefore, we look at classical, risk-neutral revenue management and directions already existing in risk-averse and robust revenue management. The resulting classification is given in Figure 2.

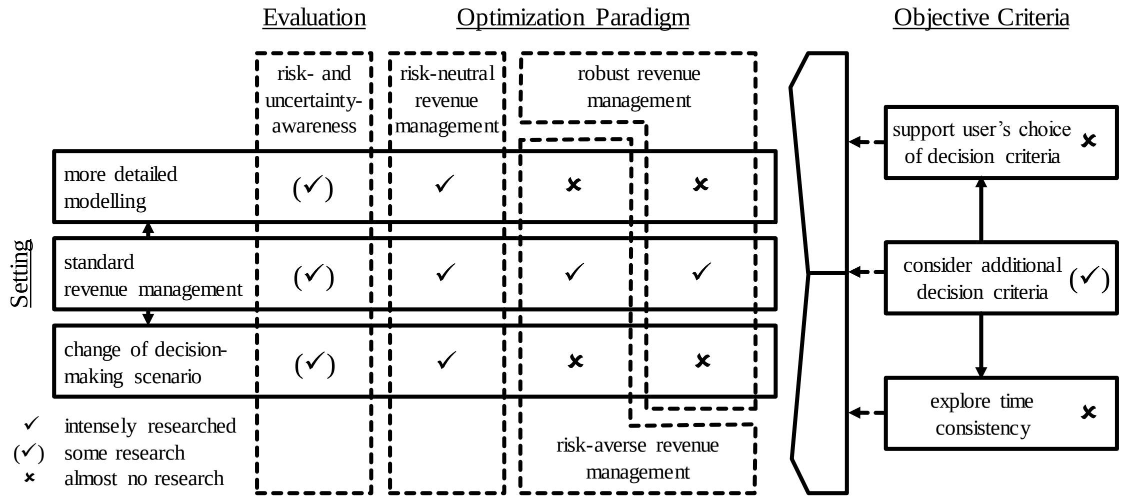

Figure 2: Potential avenues for future research

7.1 Risk- and Uncertainty-aware Revenue Management

In classical risk-neutral revenue management, most authors strictly restrict themselves to the consideration of expected values and do not expose the variability to the reader. For example, often mean values and standard deviations of simulation runs are given. However, the standard deviation refers to the mean and its purpose is to describe the accuracy of the reported mean. It does not reflect revenue variability of a single simulation run or a single sales process in practice. Only a few authors really point their readers to the fact that revenue is fluctuating and investigate the approaches considered in this regard. One such example is Koenig and Meissner (2010), who compare dynamic pricing and capacity control. Besides expected revenue, they explicitly compare the approaches' 'revenue risk' using standard deviation and the risk-measure CVaR.

Petrick et al. (2010, 2012) are uncertainty-aware as they consider biased demand forecasts, but they only report average revenue in their simulations.

Such *risk- and uncertainty-aware revenue management* is arguably neither a really independent research area nor completely new, as the above examples show. However, we put it into our list of future research to stress its importance because research is still very scarce. We feel that this is a new perspective with an extremely high relevance for practice that can be easily incorporated when simulations are used for evaluation.

7.2 Consideration of Additional Decision Criteria

In Sections 3-6, the multitude of *decision criteria* used was introduced and it became obvious that several criteria have been seldom or never considered in capacity control (chance constraints, mean-risk) and dynamic pricing (chance constraints, mean-risk, and competitive ratio). Thus, the consideration of these criteria is probably the most obvious avenue for future research. As these criteria all have been optimized in at least one setting, this should provide a starting point. Note that, Cooper and Gupta (2006) consider stochastic dominance in capacity control, but they do not use the concept as a decision criterion to evaluate different revenue distributions. They consider a risk-neutral firm and compare two markets that only differ in their demand distribution. In this setting, the authors use stochastic comparisons of the demand distributions to establish conditions that guarantee a higher expected revenue in one market.

By contrast, stochastic dominance has not been considered yet. Stochastic dominance relations are partial orders on the space of distributions (see also Lehmann (1955) as well as Hadar and Russell (1969)) and thus allow for the pairwise comparison of policies π. Stochastic dominance has an universal character with respect to utility functions. If a policy π_1 (and thus the relating revenue R_{π_1}) dominates π_2, it intuitively means that, regardless of the exact preferences, no decision maker will prefer the dominated policy π_2 and this should also be reflected in other approaches to risk-aversion. Unfortunately, it is impossible to derive simple calculation rules for stochastic dominance, due to the required pairwise comparisons (e.g., Ogryczak and Ruszczynski (2002)). Moreover, it is a very strong concept in the sense that only comparatively few pairs of distributions can be ordered and it cannot be directly used as an objective. Regarding constraints, the restriction to non-dominated policies seems desirable, but some form of non-dominance is attained by most criteria discussed in Section 3 (see also the references to stochastic dominance included there). The next subsection is closely related and discusses the topic from the decision-makers' point-of-view.

7.3 Support User's Choice of Decision Criteria

Currently, the literature on risk-aversion assumes that a *risk criterion and the degree of risk-aversion* is given. However, this is an extremely challenging issue for practitioners, especially managers with a scarce background in statistics. Thus, research could focus more on the "determination" of risk-aversion and the models sensitivity regarding this choice. Moreover, behavioral/experimental research could investigate how well managers can understand and use various criteria.

The literature on robust revenue management is more advanced in this regard, as it also gives some hints on how to select the uncertainty set.

However, the most important question underlying all this is about the real goal of the decision maker. Is it about optimizing a certain criterion? Is she really concerned about the worst case in some uncertainty set? Or just about (vaguely defined) 'good and stable' results, as numerical studies considering the expected revenue vs. standard deviation trade-off suggest? How should these goals be brought together? The well-developed literature (see, e.g., Durbach and Stewart (2012) or Greco et al. (2016) and the references therein) on multi-criteria decision making provides ample possibilities, depending on individual preferences. However, a universally adopted operational measure would be desirable, at least for evaluation purposes. While standard deviation is widely used, it has same disadvantages (see Section 3.2) and a convex combination of expected value and some coherent risk-measure might be a path to follow. In this case, optimality with regard to some surrogate criterion might lose its appeal.

While considering additional criteria is clearly relevant and necessary basic work, this all might fairly limit the benefit of just ticking the box next to another criterion.

Moreover, the different risk-averse and robust policies can be very similar regarding these measures, as a comparison of numerous types of ϕ-divergence uncertainty sets by Sierag and van der Mei (2016) for low quality demand forecasts suggests. As outlined above, progress in this regard requires some conceptual thoughts but mostly relates to evaluations and, thus, should be feasible.

7.4 Time Consistency

There is still no consensus in the literature on stochastic optimization regarding the definition of *time consistency* in a dynamic setting. Several definitions exist, all with some intuitive appeal.

- First, one approach is to consider a criterion for total revenue over the entire horizon (see, e.g. Pflug and Pichler (2016)). In revenue management, the firm is usually interested only in total revenue of one instance of the sales process (e.g. the revenue obtained from selling a flight throughout its booking horizon. Thus, almost all authors have this in mind when designing their models and evaluations. The only exceptions are time-separable utility functions and Schur et al. (2017) for CVaR.
- Second, another approach is based on risk measures which are considered time consistent if the following statement holds: If some random revenue R_1 is always (i.e. for every state of the system) riskier than another random revenue R_2 conditioned to a given time period t, then R_1 is also riskier than R_2 conditioned to the preceding time period $t + 1$. Or, more formally, knowing the value of the risk measure for all conditional distributions is sufficient to calculate its unconditional value (Artzner et al. (2007)).
- A third approach focuses on policies. According to Shapiro (2009), an optimal policy is time consistent when decisions in a given time period do not depend on future scenarios that are already known to be impossible at that point in time.

- Fourth, Rudloff et al. (2014) deem a policy “time consistent if and only if the future planned decisions are actually going to be implemented”.

Whereas the classical, risk-neutral optimization of expected values satisfies all four definitions, risk measures usually do not. Often, the first and third definition contradict. For example, the VaR considered by Koenig and Meissner (2015a) obviously satisfies the first and fourth and violates the second and third decision. Regarding CVaR, approaches satisfying the first and fourth (Gönsch et al. (2017)) as well as the third and fourth definition (Schur et al. (2017)) exist.

This issue is not even mentioned in most of the literature on revenue management. Only very few authors motivate their choice of decision criterion using some of the above definitions. Ideally, future research on time consistency with regard to revenue management should provide guidance how to navigate this issue, for example by pointing out which definitions are important and when (e.g. in what situations/industries). This should go beyond risk-measures and include risk-averse criteria as well as robust approaches because at least some definitions (number 3 and 4) seem very general. Showing whether criteria are time-consistent should be easy to accomplish, but the feasibility of developing tractable, time-consistent alternatives where necessary is an open question.

7.5 More Detailed Modelling

Compared to the classical, risk-neutral approaches, research incorporating risk-aversion and robustness is still in its infancy regarding the *modelling of the sales process*. In capacity control, for example, almost all authors still consider the basic single-leg setting without customer choice. Bridging this gap is an obvious avenue for future research and developing manageable models and heuristics that include extensions like resource networks and customer choice behavior will be necessary to render these approaches applicable to a wider range of settings in practice. The feasibility of this direction is proved by recent research like Schlosser (2015), who incorporates advertising in a dynamic pricing model.

However, there is no clear picture in this regard. On the one hand, a more detailed model itself might reduce uncertainty from misspecifications. On the other hand, an overly detailed model can be very demanding regarding the forecast and might, therefore, require robust approaches.

7.6 Change of Decision-Making Scenario

The abovementioned extensions aim at capturing the selling process in more detail. By contrast, other approaches ultimately target a *change of the decision-making* scenario, that is, for example, the information status or the selling process with its products are improved. This is probably the most promising avenue for future research, as these approaches have only been considered from the risk-neutral point of view until now, where they were already shown to have a big potential to deal with uncertain and stochastic demand. Moreover, changing the scenario should also improve the expected revenue vs. risk aversion trade-off. While optimizing decision criteria for risk-aversion and robustness will be demanding, simply moving these approaches to the stage of risk-aware revenue management and simply evaluating these criteria will already yield important insights regarding their potential to improve the risk profile of a sales process.

In the following, we briefly describe three groups of approaches. The first group improves the information status by reducing uncertainty with learning, while the second improves the reaction to occurring events through creating flexibilities on the supplier side. The third calls for the consideration of sources of variability other than demand.

7.6.1. Reducing Risk and Uncertainty through Learning

All approaches considered thus far assume an exogenously given information status that needs to be created with suitable forecasting instruments (see, e.g. Talluri and van Ryzin (2004, Chapter 9) for an overview) before the application of the optimization method. The methods described in this subsection abolish the strict separation between forecasting and optimization and improve the knowledge of demand by learning during the sales period.

Learning is intensively researched in dynamic pricing, as the almost 400 references in the literature overview by den Boer (2015) demonstrate impressively. The basic concept is the improvement in demand forecasts during the sales period by observing the past price and sales data. If the optimal price is determined only against the background of the respective current forecasts, this is called *passive learning*. *Active learning*, on the other hand, also considers the information gained by setting the current price and the value of this information for future optimization. Thus, short-term suboptimal prices, e.g. either extremely low or high, may be set. These, however, lead to a substantial improvement in the forecast and, thus, to a higher revenue in the long run. This field is methodically demanding and considers various objectives.

In contrast, learning is almost not considered in capacity control. Van Ryzin and McGill (2000), however, develop an approach that overcomes the separation of forecast and optimization, and determines protection levels directly from the historical data of a recurring sales process. The authors show that the algorithm converges to the optimal protection levels. Maglaras and Eren (2015) propose a procedure that iteratively estimates demand using maximum entropy distributions and uses the estimate for capacity control. Both approaches employ passive learning.

7.6.2. Creation of Supply-side Flexibilities

One way to create supply-side flexibility are capacity adjustments during the sales period. For example, an airline has various possibilities for short-term capacity adjustment. Ringborn and Shy (2002) consider the so-called *moving curtain* in an optimization model. Such a moving curtain keeps the total number of seats in an airplane constant, but changes the sectioning of the cabin classes. In contrast, c*onvertible seats*, which Pak et al. (2003) consider, simultaneously change the sectioning and the total number of seats. It is also possible to switch to another aircraft with a different seat configuration at short notice to better match demand and capacity. Berge and Hopperstad (1993) pursue this *flight refleeting* or *demand driven dispatch* approach. De Boer (2004) proposes EMSR-d, a variant of the popular EMSR-b heuristic. Wang and Regan (2006) focus on revenue management for two flights whose aircraft can be swapped once at a given point in time.

Furthermore, a supplier can create flexibility through *flexible products*, which offer substitution possibilities with regard to the resources used to provide a product. Gallego and Phillips (2004), as well as the influential paper of Gallego et al. (2004), introduced this concept to the literature on capacity control. Examples here are upgrades, which, for example, allow a car rental company to use higher value vehicles to serve demand for lower value vehicles (Steinhardt und Gönsch (2012)). Gönsch et al. (2014) provide additional examples of flexible products from practice as well as an extensive literature overview. Koch et al. (2017) show how to transform a setting with flexible products to the standard setting and thus render all existing network revenue management models applicable.

7.6.3. Additional Sources of Variation

In line with the classical revenue management setting, where demand is the only stochastic variable, the literature on risk-averse and robust revenue management almost solely focuses on variability caused by demand. Although demand is clearly hard to predict, there are of course additional sources of uncertainty like capacities and prices.

References

Albers S (2003) Online algorithms. A survey, *Mathematical Programming*, 97, 3-26.

Artzner P, Delbaen F, Eber J-M, Heath D (1999) Coherent measures of risk, *Mathematical Finance*, 9, 203-228.

Artzner P, Delbaen F, Eber J-M, Heath D, Ku H (2007) Coherent multiperiod risk adjusted values and Bellman's principle, *Annals of Operations Research*, 152, 5-22.

Ball M, Queyranne M (2009) Toward robust revenue management. Competitive analysis of online booking, *Operations Research*, 57, 950-963.

Barz C (2007) *Risk-averse capacity control in revenue management*, Berlin: Springer.

Barz C, Waldmann K-H (2007) Risk-sensitive capacity control in revenue management, *Mathematical Methods of Operations Research,* 65, 565-579.

Belobaba P (1989) Application of a probabilistic decision model to airline seat inventory control, *Operations Research,* 37, 183–197.

Berge M, Hopperstad C (1993) Demand driven dispatch. A method for dynamic aircraft capacity assignment, *Operations Research,* 41, 153-168.

Birbil S, Frenk J, Gromicho J, Zhang S (2009) The role of robust optimization in single-leg airline revenue management, *Management Science,* 55, 148-163.

Bitran G, Caldentey R (2003) An overview of pricing models for revenue management, *Manufacturing & Service Operations Management,* 5, 203-229.

Boda K, Filar J (2006) Time consistent dynamic risk measures, *Mathematical Methods of Operations Research*, 63, 169-186.

Chiang W, Chen J, Xu X (2007) An overview of research on revenue management, *International Journal of Revenue Management,* 1, 97–128.

Cross R (1997) *Revenue management: Hardcore tactics for market domination*, New York: Broadway Books.

Cooper W, Gupta D (2006) Stochastic comparisons in airline revenue management, *Manufacturing & Service Operations Management,* 8, 221-234.

De Boer S (2004) The impact of dynamic capacity management on airline seat inventory control*, Journal of Revenue and Pricing Management,* 2, 315-330.

Defregger F, Kuhn H (2007) Revenue management for a make-to-order company with limited inventory capacity, *OR Spectrum*, 29, 137–156.

den Boer A (2015) Dynamic pricing and learning*, Surveys in Operations Research and Management Science,* 20, 1-18.

Dentcheva D, Ruszczynski A (2006) Portfolio optimization with stochastic dominance constraints, *Journal of Banking and Finance*, 30, 433-451.

Durbach I, Stewart T (2012) Modeling uncertainty in multi-criteria decision analysis, *European Journal of Operational Research*, 223, 1-14.

Eeckhoudt L, Gollier C, Schlesinger H (1995) The risk-averse (and prudent) newsboy, *Management Science*, 41, 786-794.

Farias V, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data, *Management Science*, 59, 305-322.

Feng Y, Xiao B (1999) Maximizing revenues of perishable assets with a risk factor, *Operations Research*, 47, 337-341.

Feng Y, Xiao B (2008) A risk-sensitive model for managing perishable products, *Operations Research,* 56, 1305-1311.

Gallego G, Iyengar G, Phillips R, Dubey A (2004) Managing flexible products on a network, CORC Technical Report TR-2004-01, Columbia University.

Gallego G, Phillips R (2004) Revenue management of flexible products, *Manufacturing & Service Operations Management,* 6, 321-337.

Gallego G, Ratliff R, Shebalov S (2015) A general attraction model and sales-based linear program for network revenue management under customer choice, *Operations Research,* 63, 212-132.

Gao H, Ball M, Karaesmen I (2010) Competitive seat inventory control decisions under the regret criterion, *Journal of Revenue and Pricing Management,* 9, 49-65.

Gao H, Ball M, Karaesmen I (2016) Distribution-free methods for multi-period, single-leg booking control, *Journal of Revenue and Pricing Management,* 15, 425-453.

Geraghty M, Johnson E (1997) Revenue management saves national car rental, *Interfaces*, 27 (1), 107–127.

Gönsch J, Hassler M (2014) Optimizing the conditional value-at-risk in revenue management, *Review of Managerial Science*, 8, 495-521.

Gönsch J, Hassler M, Schur R (2017) Optimizing the conditional value-at-risk in dynamic pricing, working paper, University of Duisburg-Essen.

Gönsch J, Klein R, Neugebauer M, Steinhardt C (2013) Dynamic pricing with strategic customers, *Journal of Business Economics*, 83, 505-549.

Gönsch J, Koch S, Steinhardt C (2014) Revenue management with flexible products: The value of flexibility and its incorporation into DLP-based approaches, *International Journal of Production Economics*, 153, 280-294.

Gosavi A, Ozkaya E, Kahraman A (2007) Simulation optimization for revenue management of airlines with cancellations and overbooking, *OR Spectrum*, 29, 21-38.

Gotoh J, Takano Y (2007) Newsvendor solutions via conditional value-at-risk minimization, *European Journal of Operational Research*, 179, 80-96.

Greco S, Ehrgott M, Figueira J (2016) *Multiple criteria decision analysis: State of the art surveys*, 2nd edn., New York: Springer.

Guhlich H, Fleischmann M, Stolletz R (2015) Revenue management approach to due date quoting and scheduling in an assemble-to-order production system, *OR Spectrum*, 37, 951-982.

Hadar JW, Russell W (1969) Rules for ordering uncertain prospects, *American Economic Review*, 59, 25-34.

Haensel A, Mederer M, Schmidt H (2012) Revenue management in the car rental industry: A stochastic programming approach, *Journal of Revenue and Pricing Management*, 11, 99-108.

Hall J, Kopalle P, Pyke D (2009) Static and dynamic pricing of excess capacity in a make-to-order environment, *Production and Operations Management*, 18, 411-425.

Heching A, Gallego G, van Ryzin GJ (2002) Mark-down pricing: An empirical analysis of policies and revenue potential at one apparel retailer, *Journal of Revenue and Pricing Management*, 1, 139–160.

Hetrakul P, Cirillo C (2015) Customer heterogeneity in revenue management for railway services, *Journal of Revenue and Pricing Management,* 14, 28-49.

Hintsches A, Spengler T, Volling T, Wittek K, Priegnitz G (2010) Revenue management in make-to order manufacturing: Case study of capacity control at ThyssenKrupp VDM, *Business Research,* 3, 173–190.

Huang K, Chang K-C (2011) A model for airline seat control considering revenue uncertainty and risk, *Journal of Revenue and Pricing Management,* 11, 161-171.

Kellermann N, Cleophas C (2015) Revenue management and the railway conundrum – The consequences of reference prices in passenger railway transport practice, *Journal of Revenue and Pricing Management,* 14, 155-165.

Klein J (2000) Yield Management als Methode zur ertragsorientierten Kapazitätsnutzung bei Reiseveranstaltern—Eine empirische Untersuchung, *Tourismus-Journal*, 4, 283–307.

Klein R (2007) Network capacity control using self-adjusting bid-prices. *OR Spectrum*, 29, 39–60.

Koch S, Gönsch J, Hassler M, Klein R (2016) Practical decision rules for risk-averse revenue management using simulation-based optimization, To appear in *Journal of Revenue and Pricing Management*, 15, 468-487.

Koch S, Gönsch J, Steinhardt C (2017) Dynamic programming decomposition for choice-based revenue management with flexible products, To appear in *Transportation Science*.

Koenig M, Meissner J (2010) List pricing versus dynamic pricing: Impact on the revenue risk, *European Journal of Operational Research*, 204, 505-512.

Koenig M, Meissner J (2015a) Value-at-risk optimal policies for revenue management problems, *International Journal of Production Economics*, 166, 11-19.

Koenig M, Meissner J (2015b) Risk management policies for dynamic capacity control, *Computers & Operations Research*, 59, 104-118.

Koenig M, Meissner J (2016) Risk minimising strategies for revenue management problems with target values, *Journal of the Operational Research Society*, 67, 402-411.

Kouvelis P, Yu G (1997) *Robust discrete optimization and its applications. Series on non-convex optimization and its applications*. Amsterdam: Kluwer.

Kunnumkal S, Topaloglu H (2009) A stochastic approximation method for the single-leg revenue management problem with discrete demand distributions, *Mathematical Methods of Operations Research*, 70, 477-504.

Lai K, Ng W (2005) A stochastic approach to hotel revenue optimization, *Computers & Operations Research*, 32, 1059-1072.

Lan Y, Gao H, Ball M, Karaesmen I (2008) Revenue management with limited demand information, *Management Science*, 54, 1594-1609.

Lan Y, Ball M, Karaesmen I (2011) Regret in overbooking and fare-class allocation for single leg, *Manufacturing & Service Operations Management*, 13, 194-208.

Lan Y, Ball M, Karaesmen I, Zhang J, Liu G (2015) Analysis of seat allocation and overbooking decisions with hybrid information, *European Journal of Operational Research*, 240, 493-504.

Lancaster J (2003) The financial risk of airline revenue management, *Journal of Revenue and Pricing Management*, 2, 158-165.

Lardeux B, Goyons O, Robelin C-A (2010) Availability calculation based on robust optimization, *Journal of Revenue and Pricing Management*, 9, 313-325.

Lehmann E (1955) Ordered families of distributions, *Annals of Mathematical Statistics*, 26, 399-419.

Levin Y, McGill J, Nediak M (2008) Risk in revenue management and dynamic pricing, *Operations Research*, 56, 326-343.

Li M, Zhuang W (2009) Risk-sensitive dynamic pricing for a single perishable product, *Operations Research Letters*, 37, 327-332.

Lim A, Shanthikumar, J (2007) Relative entropy, exponential utility, and robust dynamic pricing, *Operations Research* 55, 198-214.

Lim A, Shanthikumar J, Watewai T (2008) Robust multi-product pricing, working paper, University of California Berkeley.

Maglaras C, Eren S (2015 A maximum entropy joint demand estimation and capacity control policy, *Production and Operations Management*, 24, 438-450.

Malighetti P, Paleari S, Redondi R (2009) Pricing strategies of low-cost airlines: The Ryanair case study, *Journal of Air Transport Management* 15, 195-203.

Marcus B, Anderson C (2008) Revenue management for low-cost providers, *European Journal of Operational Research,* 188, 258–272.

Markowitz H (1959) *Portfolio selection: Efficient diversification of investments*, 2nd edn., Malden: Blackwell.

Nahler, M (2013) Besonderheiten in der Kontingentierung beim Schienenverkehr, Presentation at the 10th annual meeting of the GOR working group “Pricing & Revenue Management”, Berlin, February 8th, 2013.

Noyan N, Rudolf G (2013) Optimization with multivariate conditional value-at-risk constraints, *Operations Research* 61, 990-1013.

Ogryczak W, Ruszczynski A (2002) Dual stochastic dominance and related mean-risk models, *SIAM Journal of Optimization*, 13, 60-78.

Pak K, Dekker R, Kindervater G (2003) Airline revenue management with shifting capacity, working paper, Erasmus Research Institute of Management, Erasmus University Rotterdam.

Perakis G, Roels, G (2008) Regret in the newsvendor model with partial information. *Operations Research*, 56, 188-203.

Perakis G, Roels, G (2010) Robust controls for network revenue management, *Manufacturing & Service Operations Management*,12, 56-76.

Petrick A, Gönsch J, Steinhardt C, Klein R (2010) Dynamic control mechanisms for revenue management with flexible products, *Computers & Operations Research*, 37, 2027-2039.

Petrick A, Steinhardt C, Gönsch J, Klein R (2012) Using flexible products to cope with demand uncertainty in revenue management, *OR Spectrum*, 34, 215-242.

Pflug G, Pichler A (2016) Time-inconsistent multistage stochastic programs: Martingale bounds, *European Journal of Operational Research*, 249, 155-163.

Phillips RL (2005) *Pricing and revenue optimization*, Stanford: Stanford University Press.

Prékopa A (2003) Probabilistic programming, in: Ruszczynski A, Shapiro A (eds) *Stochastic programming*, Amsterdam: Elsevier, 267-351.

Ringborn S, Shy O (2002) The ‘adjustable-curtain’ strategy – Overbooking of multiclass service, *Journal of Economics*, 77, 73-90.

Rudloff B, Street A, Valladão D (2014) Time consistency and risk averse dynamic decision models: Definition, interpretation and practical consequences, *European Journal of Operational Research*, 234, 743-750.

Rusmevichientong P, Topaloglu H (2012) Robust assortment optimization in revenue management under the multinomial logit choice model, *Operations Research*, 60, 865-882.

Schlosser R (2015) A stochastic dynamic pricing and advertising model under risk aversion, *Journal of Revenue and Pricing Management*, 14, 451-468.

Schlosser R (2016) Stochastic dynamic multi-product pricing with dynamic advertising and adoption effects, *Journal of Revenue and Pricing Management*, 15, 153-169.

Schütze J (2008) Pricing strategies for perishable products: The case of Vienna and the hotel reservation system hrs.com, *Central European Journal of Operational Research*, 16, 43–66.

Schur R, Gönsch J, Hassler M (2017) Time consistent risk averse dynamic pricing, working paper, University of Duisburg-Essen.

Seitz A, Ehm H, Akkerman R, Osman S (2016) A robust supply chain planning framework for revenue management in the semiconductor industry, *Journal of Revenue and Pricing Management* 15, 523-533.

Shapiro, A (2009) On a time consistency concept in risk averse multistage stochastic programming, *Operations Research Letters*, 37, 143–147.

Simon HA (1959) Theories of decision-making in economics and behavioral science, *American Economic Review*, 49, 253–283.

Singh P (2011) Risk in revenue management. Acknowledging risk's existence and knowing how to minimize it, *Analytics-Magazine*, 2011, 16-19.

Sierag D, van der Mei R (2016) Single-leg choice-based revenue management: A robust optimisation approach, *Journal of Revenue and Pricing Management*, 15, 454-467.

Spengler T, Rehkopf S, Volling T (2007) Revenue management in make-to-order manufacturing: An application to the iron and steel industry, *OR Spectrum*, 29, 157–171.

Steinhardt C, Gönsch J (2012) Integrated revenue management approaches for capacity control with planned upgrades, *European Journal of Operational Research*, 223, 380-391.

Talluri K, van Ryzin G (1998) An analysis of bid-price controls for network revenue management, *Management Science*, 44, 1577-1593.

Talluri K, van Ryzin G (1999) A randomized linear programming algorithm for network revenue management, *Transportation Science*, 33, 207-216.

Talluri K, van Ryzin G (2004) *The theory and practice of revenue management*, New York: Springer.

Topaloglu H (2008) A stochastic approximation method to compute bid prices in network revenue management problems, *INFORMS Journal on Computing*, 20, 596-610.

van Ryzin G, McGill J (2000) Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels, *Management Science*, 46, 760-775.

van Ryzin G, Vulcano G (2008) Simulation-based optimization of virtual nesting controls for network revenue management, *Operations Research*, 56, 865-880.

Volling T, Akyol D, Wittek K, Spengler T (2012) A two-stage bid-price control for make-to-order revenue management, *Computers & Operations Research*, 39, 1021–1032.

von Neumann J, Morgenstern O (1944) *The theory of games and economic behavior*, Princeton: Princeton University Press.

Wang X, Regan A (2006) Dynamic yield management when aircraft assignments are subject to swap, *Transportation Research Part B*, 40, 563-576.

Wang J, Xiao B (2017) A minmax regret price control model for managing perishable products with uncertain parameters, *European Journal of Operational Research*, 258, 652-663.

Weatherford L (2004) EMSR versus EMSU: Revenue or utility? *Journal of Revenue and Pricing Management*, 3, 277-284.

Xu M (2010) A price-setting newsvendor model under CVAR decision criterion with emergency procurement, *Journal of System Science and Systems Engineering*, 19, 85-104.

Xylander J (2003) *Kapazitätsmanagement bei Reiseveranstaltern: Entscheidungsmodelle zur Kontingentierung im Yield Management*, Wiesbaden: DUV.

You P (2008) An efficient computational approach for railway booking problems, *European Journal of Operational Research*, 40, 831-844.

Zhao W, Zheng Y (2000) Optimal dynamic pricing for perishable assets with nonhomogeneous demand, *Management Science*, 46, 375–388.

Zhuang W, Li M (2011) Revenue optimization of risk-averse managers with atemporal utility, *Journal of Revenue and Pricing Management*, 10, 424-437.