Time-Consistent, Risk-Averse Dynamic Pricing

Rouven Schur, Jochen Gönsch, Michael Hassler

07.01.2019

Rouven Schur

Department of Analytics & Optimization,
University of Augsburg, Universitätsstraße 16, 86159 Augsburg, Germany
E-mail: rouven.schur@wiwi.uni-augsburg.de

Jochen Gönsch (Corresponding Author)

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

Michael Hassler

Department of Analytics & Optimization,
University of Augsburg, Universitätsstraße 16, 86159 Augsburg, Germany
E-mail: michael.hassler@wiwi.uni-augsburg.de

Abstract

Many industries use dynamic pricing on an operational level to maximize revenue from selling a fixed capacity over a finite horizon. Classical risk-neutral approaches do not accommodate the risk aversion often encountered in practice. When risk aversion is considered, time-consistency becomes an important issue. In this paper, we use a dynamic coherent risk-measure to ensure that decisions are actually implemented and only depend on states that may realize in the future. In particular, we use the risk measure Conditional Value-at-Risk (CVaR), which recently became popular in areas like finance, energy or supply chain management.

A result is that the risk-averse dynamic pricing problem can be transformed to a classical, risk-neutral problem. To do so, a surprisingly simple modification of the selling probabilities suffices. Thus, all structural properties carry over. Moreover, we show that the risk-averse and the risk-neutral solution of the original problem are proportional under certain conditions, that is, their optimal decision variable and objective values are proportional, respectively. In a small numerical study, we evaluate the risk vs. revenue trade-off and compare the new approach with existing approaches from literature.

This has straightforward implications for practice. On the one hand, it shows that existing dynamic pricing algorithms and systems can be kept in place and easily incorporate risk aversion. On the other hand, our results help to understand many risk-averse decision makers who often use “conservative” estimates of selling probabilities or discount optimal prices.

Keywords: Revenue Management, Dynamic Pricing, Risk aversion, Conditional Value-at-Risk

1 Introduction

Two decades ago, product prices in the retail and service industry were rarely adjusted. The reasons for this were mainly associated with the high costs of price changes. The now ubiquitous use of the Internet as a major sales channel, digital price tags, as well as the continued development of IT software and hardware have fundamentally changed this situation (see, e.g., Talluri and van Ryzin 2004, Chap. 5.1.2 or Gönsch et al. 2013). The costs of price changes have been reduced to a minimum and sellers are now able to implement automated pricing mechanisms. This dynamic pricing aims at maximizing profits by repeatedly adjusting prices throughout the selling horizon based on the current demand and capacity situation.

In the first industries using dynamic pricing, the high repetition of events justified risk neutrality, and, thus, maximizing expected profit or revenue. This allowed to use well established tools to model dynamic decision processes in an uncertain environment, especially risk-neutral Markov decision processes/dynamic programming. However, the assumption of risk neutrality is not always appropriate, for example, when the selling process is rarely repeated (see, e.g., Feng and Xiao 1999) or a steady revenue stream is desired (see, e.g., Lancaster 2003).

1.1 Risk aversion in a static setting

Thus, the consideration of risk aversion emerged and constitutes a comparably new stream of literature. Amongst the various approaches to model risk aversion in stochastic optimization, the axiomatic approach of risk measures (respectively acceptability measures) is a prominent avenue. A risk measure ρ is a function from a set of random variables to $\mathbb{R}$. Especially the subclass of coherent risk measures has received considerable attention since the pioneering work by Artzner et al. (1999). These risk measures exhibit four properties (translation equivariance, convexity, monotonicity, and positive homogeneity) that guarantee consistency with intuition about rational risk-averse decision making. In the context of revenues, roughly speaking, translation equivariance means that adding a certain amount to the random variable equally increases the risk measure. For example, increasing future revenues for all possible states of the world by \$ 1000 would increase the risk measure by \$ 1000. Note that this implies that risk is measured in the same unit as the underlying random variables. Convexity (or subadditivity) ensures that diversifying is never bad. Monotonicity is intuitive: if revenue $R_A \geq R_B$ for all states of the world, we have $\rho(R_A) \geq \rho(R_A)$. Likewise, positive homogeneity refers to scaling: $\rho(\lambda R) = \lambda\rho(R)\ \forall\lambda \geq 0$.

This paper is based on the most prominent coherent risk measure, the Conditional Value-at-Risk (CVaR). For continuous distributions, CVaR to the probability level α is simply the conditional expectation below the α-quantile and, thus, represents the expected value of the $\alpha \cdot 100\ \%$ worst outcomes. A main feature is the fact that risk is measured in monetary units resulting in an intuitive and easily accessible interpretation of risk influenced only by the parameter α. Therefore, CVaR is easy to communicate to senior management or, more generally, to people with a scarce background in probability (see, e.g., Luciano et al. 2003 or Koenig and Meissner 2015b).

To ease the presentation, this paper largely focusses on CVaR as the risk-measure, in line with a multitude of papers using CVaR in risk-averse operations management. Examples include diverse areas such as the newsvendor problem (Chen et al. 2009, Cheng et al. 2009, Xue et al. 2015) and supply chains with returns (Hsieh and Lu 2010, Caliskan-Demirag et al. 2011). However, we also consider weighted combinations of CVaR and expected value. There, the weight allows the decision maker to express his risk aversion even finer than with CVaR alone. In the operations management literature, this combination has been used, for example, in the context of electricity production (e.g. Pousinho et al. 2012) or inventory/newsvendor problems (e.g. Ahmed et al. 2007, Gotoh and Takano 2007).

1.2 Risk aversion in a dynamic setting and time consistency

While there seems to be a general agreement in the literature on how to incorporate coherent risk measures into static models, this is not the case for the more involved dynamic models. When decisions are made dynamically and interdependently, new questions arise. How should new information be processed? How do decisions implemented and planned relate at different points in time and with a different information status? These questions are captured by the notion of time consistency. Various perspectives on time consistency have been developed in the literature so far, mainly during the last decade. They can be broadly classified into approaches focusing on risk measures and on optimal policies (see, e.g., Rudloff et al. 2014 and the references therein).

- The first approach focuses on risk-measures. A risk measure is deemed time consistent if the following statement holds: If some random revenue R_A is always (i.e. for every state of the system) riskier than another random revenue R_B conditioned to a given time period, then R_A is also riskier than R_B conditioned to the preceding time period. This approach to time consistency underlies Cheridito et al. (2006) and Ruszczyński (2010) and, thus, is equivalent to (1). It is the standard perspective for multi-period risk measures and shared by many authors, e.g., Weber (2006), Kovacevic and Pflug (2009), Riedel (2004), or the textbook by Pflug and Pichler (2014). An equivalent statement is that knowing the value of the risk measure for all conditional distributions is sufficient to calculate its unconditional value (Artzner et al. 2007).

- The second approach is standard for dynamic decision problems. Analogous to Bellman's Principle of Optimality (Bellman 1957), it states that a multi-stage stochastic decision problem is time-consistent if resolving the problem at later stages, the original solutions remain optimal. Thus, it provides the basis for deriving meaningful dynamic programming equations. For example, Bamberg and Krapp (2016), Pflug and Pichler (2016) stress this perspective on time consistency. Rudloff et al. (2014) summarize it as "a policy is time consistent if and only if the future planned decisions are actually going to be implemented". They discourage using time inconsistent models by arguing that the resulting policies are sub-optimal and propose a method to calculate the associated sub-optimality gap. Recently, Shapiro and Ugurlu (2016) showed that even with a time consistent risk-measure, multiple optimal solutions can exist, of which some might not be time-consistent. However, by recursively calculating the optimal solution, we ensure to always find a time-consistent optimal solution.
- The last approach was defined by Shapiro (2009). For optimal policies, time consistency is achieved 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.

1.3 Dynamic risk measures

For a dynamic decision problem like dynamic pricing, a dynamic risk measure is necessary. Riedel (2004) introduced the concept of dynamic coherent risk measures and describes an approach that extends static, one-period coherent risk-measures to the dynamic framework. Ruszczyński (2010) describes the concept for general Markov decision processes. Please note that we index time backwards, as usual in dynamic pricing. Given a random sequence of revenues $(R_T, R_{T-1}, \dots, R_1)$ that is adapted to the filtration $\{\emptyset, \Omega\} = \mathcal{F}_T \subset \mathcal{F}_{T-1} \subset \cdots \subset \mathcal{F}_1 \subset \mathcal{F}$, a dynamic risk measure is defined to be an $\mathcal{F}_t$-measurable sequence of conditional risk measures $\{\rho_{t,1}\}, t = T, T-1, \dots, 1$. Given a dynamic risk measure $\rho_{t,1}(R_t, R_{t-1}, \dots, R_1)$, we can derive a corresponding single-period risk measure using

$\rho_t(R_t) = \rho_{t,1}(R_t, 0, \dots, 0)$. Moreover, Ruszczyński (2010), Theorem 1, shows that any time-consistent dynamic risk measure can be constructed by nesting single-period risk measures ρ_t:

$$\rho_{t,1}(R_t, R_{t-1}, \dots, R_1) = \rho_t\Big(R_t + \rho_{t-1}\big(R_{t-1} + \dots + \rho_2\big(R_2 + \rho_1(R_1)\big) \dots\big)\Big) \tag{1}$$

and every composition of translation equivariant risk measures is time-consistent. We call (1) a nested risk measure.

The risk $\rho_{t,1}(R_t, R_{t-1}, \dots, R_1)$ is often interpreted as a certainty equivalent of the uncertain future payment stream $R_t, R_{t-1}, \dots, R_1$. In terms of costs, Ruszczyński (2010), p. 239, considers $\rho_{t,1}$ the "fair one-time $\mathcal{F}_t$-measurable charge we would be willing to incur at time t, instead of the sequence of random future costs $R_t, R_{t-1}, \dots, R_1$. In a similar way, Riedel (2004) interprets $\rho_{t,1}(R_t, R_{t-1}, \dots, R_1)$ as the minimum amount of money one has to add to the position yielding $R_t, R_{t-1}, \dots, R_1$ to make it acceptable.

As we use identical single-period risk measures $\rho_t = CVaR_\alpha\ \forall t$, equation (1) is directly amenable to a dynamic programming formulation and can be optimized stage-wise (see Ruszczyński and Shapiro 2006). This ensures time consistency in the sense of the three popular and intuitively desirable properties described in Section 1.2.

In the context of minimizing costs, Shapiro et al. (2013) motivate nested CVaR as an approach that penalizes extreme costs above a certain upper limit. This limit is not fixed but adapted, i.e. depends on the conditional distribution of future total costs given the current history. Now, the optimal decisions obtained with the $(1-\alpha)$-quantile of the conditional distribution as the upper limit and a penalty factor of $1/\alpha$ equal those for $CVaR_\alpha$ at that point in time. Rudloff et al. (2014) provide a different economic interpretation. They prove that "the objective function is the certainty equivalent w.r.t. the time consistent dynamic utility generated by one-period preference functionals". They illustrate this with an investor who is willing to forgo the uncertain future value of his portfolio and sell it now for a certain amount of money, the certainty equivalent. Thus, they can interpret the optimal certainty equivalent as the portfolio value. In dynamic pricing, the certainty equivalent is the value the firm assigns to continuing the sales process throughout the selling horizon based on the current demand and capacity situation. Or, in other words, it is the deterministic amount of money the firm needs to be offered to stop selling. Street (2010) analyzes CVaR's preference functional in detail.

Nonetheless, one may wish for a more intuitive interpretation. For example, it may be desirable to optimize a well-known risk measure over the entire selling horizon by repeatedly applying it in the stage-wise optimization. However, this is unfortunately not possible for coherent risk measures other than the expectation and the essential infimum. One has to accept one of the following three (see, e.g., Pflug and Pichler (2014), Chapter 5) downsides: (1) change the optimality criterion at later stages, and, thus, the actual criterion used may depend on the random variable. This approach is time-inconsistent according to Shapiro (2009) and followed by Gönsch et al. (2018) to optimize CVaR over the entire selling horizon by dynamically adjusting the risk level α. (2) when resolving the problem at later stages the original solutions are no longer optimal, or work with the changed decisions right from the beginning, although they are suboptimal for the unconditional problem. Again, this is by definition not time-consistent and can be realized using the algorithms of Gönsch et al. (2018) if the policy is reevaluated later with the initial risk level. Finally, (3) we can construct a time-consistent dynamic risk measure via nesting as in equation (1), albeit the measure we obtain is difficult to interpret. This is what we do in this paper.

Likewise, a nested objective function is widely used in the literature to ensure time-consistency and to incorporate risk measures into dynamic programs. For example, Shapiro (2011), Shapiro et al. (2013), and Philpott et al. (2013) focus on solution methods with nested CVaR or nested coherent measures of risk in general in the context of stochastic dual dynamic programming. Collado et al. (2012) consider decomposition of multistage stochastic programming problems. Nested risk measures are also common in operations management applications (see, e.g., Ahmed et al. 2007 for a multi-period newsvendor problem or Philpott and Matos 2012 as well as Maceira et al. 2015 for hydrothermal scheduling in New Zealand and Brasil, respectively).

1.4 Contribution

The basis for our paper is the dual representation theorem established in Artzner et al. (1999). It states that coherent risk measures can be computed by an expectation with an adjusted probability distribution. This approach is also referred to as change of probability distribution or change de numéraire. However, determining the adjusted probability distribution for calculation of CVaR at each stage usually requires the solution of an optimization problem. By contrast, we show that the problem structure is largely retained in dynamic pricing because the adjusted distribution is static in the sense that it depends only on the original distribution and the level of risk aversion and, thus, can be computed offline in advance.

A result of this paper is that the risk-averse dynamic pricing problem can be transformed into a classical, risk-neutral dynamic pricing problem with modified selling probabilities. This transformed problem has the same objective value and the same optimal solution, although a straightforward transformation for a boundary solution is necessary. Thus, all structural properties known from risk-neutral dynamic pricing carry over and new, risk-related properties are shown. Moreover, we show that the risk-averse and the risk-neutral solution of the original problem are proportional under certain conditions, that is, their optimal decision variable and objective values are proportional, respectively.

1.5 Outline

The remainder of the paper is structured as follows. Relevant literature is reviewed in Section 2. In Section 3, the risk-neutral and risk-averse dynamic pricing problems are formally stated. Based on this, structural results are presented in Section 4. Section 5 contains a numerical study that analyses the risk vs. revenue trade-off and compares the new approach to existing ones from literature. Section 6 concludes with managerial implications and an outlook on possible future research. The proofs are contained in the appendix.

2 Literature review

In the following, we review related literature. We start with methodological work on time consistency, then we discuss prior research on risk-averse dynamic pricing.

2.1 Dynamic risk measures, time consistency and multistage optimization

The literature on solving multistage stochastic programs with coherent risk measures is relevant and, especially in the area of scenario-based stochastic dual dynamic programming (SDDP), related results exist (see, e.g., Shapiro et al. 2009 (Chapter 6), Shapiro 2011, and Shapiro et al. 2013). Shapiro (2011) builds on a formulation of CVaR described in Pflug (2000) and Rockafellar and Uryasev (2000) which calculates CVaR at each stage by optimizing an additional decision variable that corresponds to VaR in the optimum. Philpott and Matos (2012) demonstrate that although this approach uses one additional decision variable per stage, it solves large-scale stochastic programming problems. In their Sample Average Approximation problem, Shapiro et al. (2013) circumvent the additional decision variable and calculate the VaR at each stage by ordering the scenarios. Philpott et al. (2013) build on the idea of ordering the scenarios. However, they use CVaR's dual representation due to Artzner et al. (1999) and order events according to their objective value to calculate the adapted probability measure at each stage. This idea is applicable to any coherent risk measure and allows optimizing an expected value. Our results are in line with this. However, there is an important difference between the settings. Philpott et al. (2013) separately optimize the stochastic outcomes (and, thus, the expected value) because the probability distribution is exogenous. By contrast, we consider an endogenous probability distribution (the selling probability depends on the price) and optimize by affecting the outcomes as well as the probabilities. In addition, our transformation of the probability distribution is static in the sense that it is carried out offline in advance and we obtain the same, albeit risk-neutral problem.

2.2 Risk-averse dynamic pricing

The field of dynamic pricing emerged about 30 years ago from the study of intertemporal price discrimination (see, e.g., Stokey 1979, Landsberger and Meilijson 1985, and Wilson 1988). The onset of modern-day research in this field can be attributed to the seminal paper by Gallego and van Ryzin (1994), who considered optimal dynamic pricing of a single product with stochastic demand over a finite selling horizon. Since the publication of this paper, research on dynamic pricing has increased significantly with the publication of hundreds of follow-up papers. Several review articles (e.g., Bitran and Caldentey 2003, Chiang et al. 2007, and, with a special focus, Gönsch et al. 2013 and den Boer 2015) as well as textbooks (e.g. Talluri and van Ryzin 2004 (Chapter 5) and Phillips 2005 (Chapter 10)) have structured and summarized this research.

Whereas there is a large body of literature on risk-neutral dynamic pricing, the consideration of risk aversion constitutes a rather new field. Relevant literature is summarized in a survey on risk aversion in revenue management and dynamic pricing by Gönsch (2017). We review the most relevant papers here. Feng and Xiao (1999) were the first to introduce risk aversion in a dynamic pricing framework. Choosing from a pair of pre-determined prices, the decision maker maximizes expected revenue with an additional penalty term for revenue variance that takes business risk into account. Levin et al. (2008) follow a target-criterion of the satisficing type, i.e., a fixed minimum revenue has to be attained with at least a given probability. Several authors capture the decision maker's risk aversion via utility functions based on the expected utility theory of von Neumann and Morgenstern (1944). Lim and Shanthikumar (2007) connect risk-averse dynamic pricing to robust optimization by showing the equivalence of risk-averse, single-product dynamic pricing with an exponential utility function and robust dynamic pricing which takes demand model errors into account. When using additive general utility and atemporal exponential utility functions, the well-known monotonicities from risk-neutral dynamic pricing are preserved under risk aversion, which was shown by Li and Zhuang (2009). They also show that the optimal price decreases in risk aversion, i.e., more risk-averse decision makers tend to set lower prices. Schlosser (2015) includes the decision about the advertising intensity, which, just as the sales price, influences sales and derives optimal closed-form policies for exponential utility functions.

To the best of our knowledge, Gönsch et al. (2018) is the only paper that uses coherent risk measures as a target criterion in a dynamic pricing framework. The authors recursively maximize CVaR of overall revenue over the whole time horizon, following Pflug and Pichler (2016). Although intuitive, this target criterion is not time consistent in the sense of Section 1.2 and the resulting decisions can depend on scenarios that already became impossible. In contrast, we ensure time consistency and work with a nested objective function resulting from the recursive application of one-period CVaR following Shapiro (2009). In addition to that, CVaR has been used by Koenig and Meissner (2010) to evaluate (but not to optimize) dynamic pricing policies.

Finally, note that risk aversion has also been considered in capacity control, where product availability instead of price is the decision variable (see, e.g., Barz and Waldmann 2007, Gönsch and Hassler 2014, Huang and Chang 2011, Koch et al. 2016, Koenig and Meissner 2016 as well as the survey by Gönsch 2017 and the references therein).

3 Problem definition

In this section, we formally define the problem considered. We first introduce the setting and notation and restate the risk-neutral dynamic pricing problem. Then, CVaR is formally introduced. Based on this, we state the risk-averse dynamic pricing problem with nested CVaR.

3.1 Setting and notation

We consider the standard setting of dynamic pricing. A firm optimizes its revenue from selling an initial stock of $C \in \mathbb{N}$ units of a single product during a selling horizon of length $T \in \mathbb{N}$ through price variations. The selling horizon is discretized and indexed backwards in time, i.e., periods T and 0 mark the beginning and the end of the horizon. Any capacity remaining afterwards is worthless. We assume that exactly one customer arrives in each time period. More specifically, at the beginning of each time period $t \in \{T, \dots, 1\}$, the firm knows the remaining capacity c and sets a price p. To simplify notation, we assume that exactly one customer arrives after this. Her willingness-to-pay (WTP) X_t is a continuous random variable with probability density function $f_t(x_t)$ and cumulative distribution function $F_t(x_t)$ that is independent between periods. She buys the product if and only if the price p does not exceed her WTP X_t. To ease notation, we write $d_t(p) = 1 - F_t(p)$ to denote the probability that a sale takes place at price p. If we had assumed that at most one customer arrives in a period, $d_t(p)$ would contain the arrival probability in addition to $1 - F_t(p)$. This would be technically equivalent as there is still at most one sale with probability $d_t(p)$. Obviously, it is not necessary to set prices above the maximum WTP and we can restrict the set of allowable prices to $\mathcal{P}_t = [0, p_t^\infty]$, where the null price p_t^∞ denotes the smallest price (possibly $+\infty$) with $\lim_{p \to p_t^\infty} d_t(p) \cdot p = 0$. The existence of a null price is a standard assumption in dynamic pricing since this allows to model the out-of-stock condition (see, e.g., Gallego and van Ryzin 1994).

Moreover, we assume that standard assumptions from literature hold (see, e.g., Ziya et al. 2004 or Talluri and van Ryzin 2004 (p. 317)). In particular, Ziya et al. (2004) structure widely used assumptions regarding demand functions, albeit in a static setting, on which we will focus in the following. More precisely, we assume that the selling probability is twice continuously differentiable, strictly decreases in price ($d'_t(p) < 0 \ \forall\, t \in \{1, 2, \dots, T\}, p \in (0, p_t^\infty)$) and that any one of the following two standard assumptions holds:

- $-\frac{d_t''(p)}{d_t'(p)} < -2\frac{d_t'(p)}{d_t(p)} \qquad \forall\, t \in \{1, 2, \dots, T\}, p \in (0, p_t^\infty)$
- $-\frac{d_t''(p)}{d_t'(p)} < -\frac{d_t'(p)}{d_t(p)} + \frac{1}{p} \qquad \forall\, t \in \{1, 2, \dots, T\}, p \in (0, p_t^\infty)$

Our first assumption is equivalent to their assumption A1/C1 and has the interpretation that the revenue function is strictly concave in demand. Our second assumption is equivalent to their assumption A3/C3 and means that the WTP distribution $F_t(x_t)$ has a strictly increasing generalized failure rate. They also show that the assumptions are not equivalent. Hempenius (1970) considers a profit function that involves a cost term and shows that any of these assumptions ensures that the solution to the first-order condition of the profit function with respect to price is an optimal solution if the cost function is convex in demand. The link from the static setting to our dynamic setting is as follows. Even though there is no direct cost associated with selling a unit in the dynamic setting, there is an indirect cost because capacity may be scarce and each unit sold today cannot be sold in the future. This indirect cost is usually referred to as opportunity cost and will be formally defined in Section 4.1. Regarding the determination of the optimal price at a given point in time, the only difference between the static and dynamic setting is that in the dynamic setting there is an (opportunity) cost, which is linear in demand. Thus, the assumptions ensure that the problem is well behaved, that is, it suffices to consider the f.o.c to obtain an optimal solution.

These regularity conditions are met by many distributions, e.g. the uniform and the exponential distribution. The first assumption is widely-used in the literature on (dynamic) pricing (e.g. Feichtinger and Hartl 1985, Li 1988, Gallego and van Ryzin 1994, Paschalidis and Tsitsiklis 2000, Bitran and Mondschein 1997, Cachon and Lariviere 2001). The second assumption is used, for example, by Lariviere and Porteus (2001).

To reduce notation, we usually omit the domain of variables and parameters whenever the corresponding standard domain is meant. In particular, where not otherwise stated it holds that $t \in \{1, 2, \ldots, T\}$, $c \in \{1, 2, \ldots, C\}$, $\alpha \in (0,1]$ and $p \in \mathcal{P}_t$.

3.2 Risk-neutral dynamic pricing

In traditional dynamic pricing, a risk-neutral firm maximizes the total expected revenue over the remaining periods t of the selling horizon with a stock of c units to sell. This is captured by the following Bellman equation:

$$V_t^1(c) = \max_{p \in \mathcal{P}_t} \mathbb{E}\left[1_{\{X_t \geq p\}} \cdot p + V_{t-1}^1\left(c - 1_{\{X_t \geq p\}}\right)\right] \tag{2}$$

where $1_{\{x\}}$ is the indicator function that equals one if and only if x is true. Here, $V_t^1(c)$ denotes the optimal expected revenue-to-go from period t onwards. The price p is set to maximize the expected sum of immediate revenue (obtained if the current customer's WTP X_t is at least p) and the revenue to go from the next period onwards with the remaining capacity. Obviously, the expectation captures two possible events: A sale occurs with probability $d_t(p)$ and the firm immediately obtains a revenue of p and additionally expects a revenue of $V_{t-1}^1(c-1)$ with a reduced stock of $c-1$ units from the next period onwards. No sale occurs with probability $1 - d_t(p)$. In this case, the firm expects a revenue of $V_{t-1}^1(c)$ from stock c. This leads to the more common formulation

$$V_t^1(c) = \max_{p \in \mathcal{P}_t} d_t(p) \cdot \left(p + V_{t-1}^1(c-1)\right) + \left(1 - d_t(p)\right) \cdot V_{t-1}^1(c) \tag{3}$$

We denote the optimal risk-neutral price selected in a state (t, c) by $p_{t,c}^1$. In (2) and (3), boundary conditions ensure termination of the recursion and the sale of at most C items: $V_0^1(c) = 0$ for $c \geq 0$ and $V_t^1(c) = -\infty$ for $c < 0$.

3.3 Representations of CVaR

Given a probability level $\alpha \in (0,1]$ and a random variable R denoting a profit with distribution function $F_R(y)$, CVaR is given by its well-known dual representation (Artzner et al. 1999):

$$\mathrm{CVaR}_\alpha(R) = \inf_Z \{\mathbb{E}[RZ] : \mathbb{E}[Z] = 1,\ 0 \leq Z \leq 1/\alpha\} \tag{4}$$

In (4), the infimum is over all nonnegative random variables $Z \geq 0$ with expectation $\mathbb{E}[Z] = 1$, which satisfy the additional truncation constraint $Z \leq 1/\alpha$. For continuous distributions, CVaR can be intuitively defined using the Value-at-Risk (VaR), which is simply the α-quantile ($\mathrm{VaR}_\alpha(R) = F_R^{-1}(\alpha)$). Then, CVaR_α equals the expectation below VaR_α or the α-quantile: $\mathrm{CVaR}_\alpha(R) = \mathbb{E}[R : R \leq F_R^{-1}(\alpha)]$.

3.4 Nested CVaR in dynamic pricing

A dynamic pricing formulation optimizing nested CVaR can now be obtained as follows. On an intuitive level, the expectation in (2) is simply replaced with CVaR. More formally, we maximize (1) with $\rho_t = CVaR_\alpha\ \forall t$ and we obtain the following dynamic programming equation (see, e.g., Ruszczyński and Shapiro 2006 for a detailed derivation of dynamic programming equations for nested risk measures):

$$V_t^\alpha(c) = \max_{p \in \mathcal{P}_t} \mathrm{CVaR}_\alpha \left(1_{\{X_t \geq p\}} \cdot p + V_{t-1}^\alpha\left(c - 1_{\{X_t \geq p\}}\right)\right) \tag{5}$$

Analogous to (2), the expectation inherent in CVaR (see equation (4)) is over two events: a sale with probability $d_t(p)$ and no sale with probability $1 - d_t(p)$. Thus, using (4), equation (5) can be rewritten as follows:

$$V_t^\alpha(c) = \max_{p \in \mathcal{P}_t} \inf_Z \{\mathbb{E}[RZ] : \mathbb{E}[Z] = 1,\ 0 \leq Z \leq 1/\alpha\} \text{ with } R = 1_{\{X_t \geq p\}} \cdot p + V_{t-1}^\alpha\left(c - 1_{\{X_t \geq p\}}\right) \tag{6}$$

Since the sample space consists of only two elements indicating if a sale occurs or not, denote by $z_{t-1,c-1}$ and $z_{t-1,c}$ the respective value of the random variable Z. Then the optimization problem in equation (5) can explicitly be written as follows:

$$V_t^\alpha(c) = \max_{p\in\mathcal{P}_t} \min_{\substack{z_{t-1,c-1},\\ z_{t-1,c}}} \big(1-d_t(p)\big)\cdot z_{t-1,c}\cdot V_{t-1}^\alpha(c) + d_t(p)\cdot z_{t-1,c-1}\cdot\big(p+V_{t-1}^\alpha(c-1)\big) \quad (7)$$

subject to

$$1 = \big(1-d_t(p)\big)\cdot z_{t-1,c} + d_t(p)\cdot z_{t-1,c-1}$$

$$0 \le z_{t-1,c} \le \frac{1}{\alpha}$$

$$0 \le z_{t-1,c-1} \le \frac{1}{\alpha}$$

with the boundary conditions $V_t^\alpha(c) = -\infty$ for $c<0$, $V_0^\alpha(c)=0$ for $c\ge 0$ applying to (5), (6), and (7). Here, $z_{t-1,c-1}$ and $z_{t-1,c}$ are the values of the random variable Z (see equation (4)) in case of the two events sale and no sale, respectively. The optimal price selected in state (t,c) is denoted by $p_{t,c}^\alpha$.

4 Structural results

This section presents several structural results. We first show that the risk-averse dynamic pricing problem (7) described in Section 3.4 can be transformed to an equivalent, risk-neutral standard dynamic pricing problem with modified selling probability, but identical objective value and solution (Section 4.1) and state well-known and new monotonicities (Section 4.2). In Section 4.3, we focus a combination of CVaR and expected value. We then focus on time-homogeneous demand and show that the optimal prices and objective values of a risk-averse and a risk-neutral decision maker are proportional if the distribution of WTP satisfies a certain condition (Section 4.4). Finally, we illustrate this using uniformly ($U[0,1]$) distributed WTPs (Section 4.5).

4.1 Transformation to risk-neutral dynamic pricing problem

In this section, we consider arbitrary demand. The only requirement is that $d_t(p)$ satisfies the regularity conditions (see Section 3.1) in every period t.

Now consider a standard, risk-neutral dynamic pricing problem with the transformed selling probability $\tilde{d}_t^\alpha(p) = 1-\big(1-d_t(p)\big)/\alpha\ \forall p\in\tilde{\mathcal{P}}_t$. The corresponding null price $p_t^{\alpha,\infty}$ is chosen such that $d_t(p_t^{\alpha,\infty}) = 1-\alpha$. Accordingly, we have $\tilde{p}_{t,c}^\alpha \in \tilde{\mathcal{P}}_t = \big[0, p_t^{\alpha,\infty}\big]$, $\tilde{d}_t^\alpha(p)\ge 0$ and the Bellman equation

$$\tilde{V}_t^\alpha(c) = \max_{p\in\tilde{\mathcal{P}}_t} \tilde{d}_t^\alpha(p) \cdot \left(p + \tilde{V}_{t-1}^\alpha(c-1)\right) + \left(1 - \tilde{d}_t^\alpha(p)\right) \cdot \tilde{V}_{t-1}^\alpha(c) \tag{8}$$

with the boundary conditions $\tilde{V}_0^\alpha(c) = 0$ for $c \geq 0$ and $\tilde{V}_t^\alpha(c) = -\infty$ for $c < 0$. We denote the optimal risk-neutral price selected in a state (t, c) by $\tilde{p}_{t,c}^\alpha$. Now, the following theorem presents our main result and relates the risk-neutral problem (8) to the risk-averse dynamic pricing problem (7).

Theorem 1 *For the risk-averse dynamic pricing problem (7) with objective $V_t^\alpha(c)$, a corresponding risk-neutral dynamic pricing problem (8) with objective $\tilde{V}_t^\alpha(c)$, transformed selling probability $\tilde{d}_t^\alpha(p)$, and null price $p_t^{\alpha,\infty}$ exists, and the following holds:*

a) The objective values of both problems are equal: $V_t^\alpha(c) = \tilde{V}_t^\alpha(c)\ \forall\ t, c$.

b) The optimal risk-averse price $p_{t,c}^\alpha$ in all states (t, c) is given by the transformed risk-neutral problem (and vice versa):

$$p_{t,c}^\alpha = \begin{cases} \tilde{p}_{t,c}^\alpha, & \tilde{p}_{t,c}^\alpha \in \left[0, p_t^{\alpha,\infty}\right) \\ p_t^\infty, & \tilde{p}_{t,c}^\alpha = p_t^{\alpha,\infty} \end{cases} \quad \forall\ t, c$$

Proof: The proof of this theorem is given in Appendix A.1.

Lemma 1 *The transformation of the selling probability described in Theorem 1 preserves all standard assumptions (see Section 3.1). Thus, the corresponding optimal price $p_{t,c}^\alpha$ is well defined.*

Proof: The proof of this lemma is given in Appendix A.1.

Lemma 2 *Let $\Delta_{t,c}^\alpha = V_{t-1}^\alpha(c) - V_{t-1}^\alpha(c-1)$. If $d_t\left(\Delta_{t,c}^\alpha\right) > 1 - \alpha$, the optimal solution $p_{t,c}^\alpha$ of the risk-averse problem is an element of the non-empty interval $\left(\Delta_{t,c}^\alpha, p_t^{\alpha,\infty}\right)$. Otherwise, $p_{t,c}^\alpha = p_t^\infty$.*

Proof: The proof of this lemma is given in Appendix A.2.

Theorem 1 can be explained as follows using Lemma 2. CVaR is the conditional expected value taking the α-worst outcomes into account. Moreover, selling (as opposed to not selling) is always the better outcome, given that a price above opportunity costs is set. This is the case because by contradiction, we see that setting a price lower than the opportunity cost makes no sense because then, selling would be the undesired event. A higher price (up to the opportunity cost) would simultaneously decrease the probability of this undesired event and improve its outcome. Thus, setting a price lower than the opportunity costs cannot be optimal. Now, the theorem can be illustrated by distinguishing two cases according to Lemma 2. Remember that $d_t(p)$ decreases in p.

In the first case, we have $d_t\left(\Delta_{t,c}^{\alpha}\right) > 1 - \alpha$. Thus, prices exceeding opportunity cost whose selling probability exceeds $1 - \alpha$ exist. Hence, the event of selling can be partly included in CVaR's expectation stretching over the α-worst outcomes (Figure 1). To see this, first, consider the upper part of Figure 1. Blocks visualize the two possible events, selling and not selling. The height of a block visualizes the value the firm obtains if the corresponding event occurs. The width of a block denotes the event's probability. The left block stands for not selling with value $V_{t-1}^{\alpha}(c)$ and probability $1 - d_t\left(p_{t,c}^{\alpha}\right)$. The right block represents the event of selling with a probability (width) of $d_t\left(p_{t,c}^{\alpha}\right)$. It consists of two stacked rectangles because its value is the sum of the price $p_{t,c}^{\alpha}$ obtained for the unit sold and the value $V_{t-1}^{\alpha}(c-1)$ obtained in the future with one unit of capacity less. Their horizontal order represents the structure of the solution of CVaR's dual representation, that is, the values of $z_{t-1,c}$ and $z_{t-1,c-1}$ (for not selling and selling, respectively) the inner maximization in (7) that calculates CVaR chooses. Now, if a price exceeding opportunity cost is chosen, the value for selling is higher than the value for not selling. Thus, it is obvious from the objective (7) and the equality constraint that in an optimal solution of the calculation of CVaR we have $z_{t-1,c-1} > 0$ if and only if $z_{t-1,c} = 1/\alpha$. We now interpret $\alpha \cdot z_{t-1,\cdot}$ as the proportion of an event's inclusion in CVaR's expectation, for example, $z_{t-1,c-1} = 0$ means that selling does not contribute, $z_{t-1,c-1} = \frac{1}{2\alpha}$ means it is half included, and $z_{t-1,c-1} = \frac{1}{\alpha}$ means that it is fully included. Additionally denoting the level α on the vertical axis now allows us to graphically solve CVaR's minimization. To do this, the $z_{t-1,\cdot}$ must simply reflect the proportion of a block covered by α. In the figure, not selling is fully included ($z_{t-1,c} = \frac{1}{\alpha}$), whereas selling is only partly included with one third ($z_{t-1,c-1} = \frac{1}{3\alpha}$).

Second, consider the lower part of Figure 1. By changing $d_t(p)$ to $\tilde{d}_t^{\alpha}(p)$, we get the conditional probability function given the α-worst outcomes. We then use this probability function to technically calculate an expected value as in the risk-neutral problem, which enables the transformation of the risk-averse to a risk-neutral dynamic pricing problem. From Lemma 2, we know that the optimal price is an inner solution, i.e. $p_{t,c}^{\alpha} \in \left(\Delta_{t,c}^{\alpha}, p_t^{\alpha,\infty}\right)$.

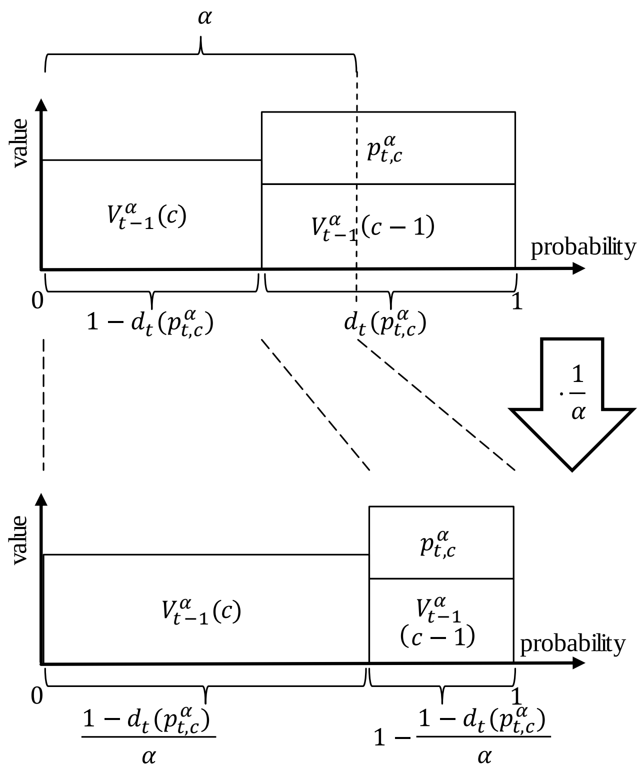

Figure 1: Illustration of case 1

In the second case, we have $d_t(\Delta_{t,c}^{\alpha}) \leq 1-\alpha$, and the time period is skipped in the sense that the firm does not take a sale into account and waits for better time periods with higher selling probabilities in the future. Now, all prices exceeding opportunity cost have a selling probability less than $1-\alpha$. Thus, the event of selling is not considered in CVaR's expectation because the α-worst outcomes include only the event of not selling with its probability exceeding α (Figure 2) and value $V_{t-1}^{\alpha}(c)$. For $p = \Delta_{t,c}^{\alpha}$, both outcomes have the same value: $V_{t-1}^{\alpha}(c) = \Delta_{t,c}^{\alpha} + V_{t-1}^{\alpha}(c-1)$. Moreover, this value is also obtained for the price with selling probability zero ($p = p_t^{\infty}$). Thus, we have $\mathrm{CVaR}_{\alpha} = V_{t-1}^{\alpha}(c)\ \forall\, p$ with $\Delta_{t,c}^{\alpha} \leq p \leq p_t^{\infty}$ (or $p = p_t^{\infty}$ if the interval $[\Delta_{t,c}^{\alpha}, p_t^{\infty}]$ is empty) and all these prices are optimal. In Theorem 1b) we set $p_{t,c}^{\alpha} = p_t^{\infty}$—the price at which the selling probability becomes zero—because the interval may be empty and because p_t^{∞} is independent of $\Delta_{t,c}^{\alpha}$ and, thus, can be easily calculated in advance. Note that it is not possible to choose the price at which the selling probability in the transformed, risk-neutral problem becomes zero ($p_{t,c}^{\alpha} = p_t^{\alpha,\infty}$), because this can be below $\Delta_{t,c}^{\alpha}$ and selling would then be strictly worse than not selling, leading to a lower CVaR than pricing at p_t^{∞}.

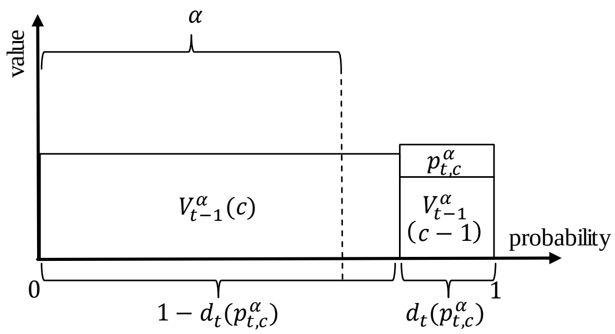

Figure 2: Illustration of case 2

Lemma 3 *With time-homogeneous demand,* $d_t\left(\Delta_{t,c}^{\alpha}\right) > 1-\alpha$ *holds for all* c, t.

Proof: The proof of this lemma is given in Appendix A.3.

Together with Lemma 2, Lemma 3 implies that with time-homogeneous demand, a price with purchase probability zero is never optimal and the optimal price lies in the interval $\left(\Delta_{t,c}^{\alpha}, p_t^{\alpha,\infty}\right)$. Otherwise, a counterexample can easily be derived by using different domains for the distributions of reservation prices such that Case 2 of Lemma 2 applies. Thus, there are no worthless time periods where selling is not considered.

4.2 Monotonicities

Theorem 1 shows how to transform the risk-averse dynamic pricing problem into a classical risk-neutral one. This transformation has a huge impact on the solvability of the problem. Over the years, the literature on dynamic pricing has tackled many issues such as stating properties of the optimal objective values as well as the optimal solutions and developing (approximate) solution methods. Using the results of Theorem 1, the existing theory and solution methods for risk-neutral dynamic pricing can immediately be applied to risk-averse dynamic pricing. For example, the well-known monotonicities regarding time and capacity carry over:

Proposition 1 *The marginal value of capacity is increasing in* t *and decreasing in* c*, i.e.,* $\Delta_{t,c}^{\alpha} \geq \Delta_{t-1,c}^{\alpha}$ *and* $\Delta_{t,c}^{\alpha} \leq \Delta_{t,c-1}^{\alpha}$ *for all* c, t. *Consequently, the optimal price* $p_{t,c}^{\alpha}$ *is increasing in* t *and decreasing in* c*, i.e.,* $p_{t,c}^{\alpha} \geq p_{t-1,c}^{\alpha}$ *and* $p_{t,c}^{\alpha} \leq p_{t,c-1}^{\alpha}$ *for all* c, t.

Proof Proposition 1 directly follows from Theorem 1, as the results are well-known to hold in risk-neutral dynamic pricing (see, e.g., Talluri and van Ryzin 2004, pp. 188 and 203–204).

Finally, a monotonicity regarding the new parameter α arises.

Proposition 2 *The optimal objective value $V_t^\alpha(c)$ increases in the risk level α.*

Proof The proof can be outlined as follows. Consider the transformed problems for two risk levels $\alpha' < \alpha''$. From the assumptions in Section 3.1 follows that the optimal policy defined by the optimal prices $p_{t,c}^{\alpha\prime}\ \forall(t,c)$ for α' implies a selling rate $d_t^{\alpha'}\left(p_{t,c}^{\alpha\prime}\right)$ for each state (t,c). We now construct a policy defined by prices $p_{t,c}^{\alpha\prime\prime}\ \forall(t,c)$ for α'' such that the selling rates in the problem with α'' match the rates in the former problem, i.e. $\tilde{d}_t^{\alpha'}\left(p_{t,c}^{\alpha\prime}\right) = \tilde{d}_t^{\alpha''}\left(p_{t,c}^{\alpha\prime\prime}\right)\ \forall(t,c)$. Note that $p_{t,c}^{\alpha\prime\prime}$ is not required to be optimal but is only chosen to mirror the given selling rates $\tilde{d}_t^{\alpha'}\left(p_{t,c}^{\alpha\prime}\right)$. With the definition of $d_t(p)$ and continuously distributed willingness-to-pays, there is a unique set of prices $p_{t,c}^{\alpha\prime\prime}$ that fulfills this requirement. Moreover, it holds that $p_{t,c}^{\alpha\prime} < p_{t,c}^{\alpha\prime\prime}\ \forall(t,c)$ as $\tilde{d}_t^{\alpha'}\left(p_{t,c}^{\alpha\prime}\right) = \tilde{d}_t^{\alpha''}\left(p_{t,c}^{\alpha\prime\prime}\right) \Leftrightarrow \left(1 - d_t\left(p_{t,c}^{\alpha\prime}\right)\right)/\alpha' = \left(1 - d_t\left(p_{t,c}^{\alpha\prime\prime}\right)\right)/\alpha''$ and with $\alpha' < \alpha''$ it follows that $d_t\left(p_{t,c}^{\alpha\prime\prime}\right) < d_t\left(p_{t,c}^{\alpha\prime}\right)$. This obviously requires higher prices for α''. The important issue is now that both problems are equivalent regarding the (stochastic) evolution of the selling process. That is, the probability of each possible evolution of the selling process is the same, while revenues are strictly higher for α''. Thus, the objective value resulting from the optimal policy of the lower risk level α' is less than the objective value resulting from the (not even necessarily optimal) policy we constructed for the higher risk level α'', which completes the proof of Proposition 2.

The meaning of Proposition 2 is straightforward. As CvaR_α monotonically increases in α, the objective value, which is a recursive calculation of CvaR_α, also increases in α.

4.3 Combination of expected value and CVaR

To ease notation, we focused on CVaR only so far. However, it is possible to adapt Theorem 1 to the following coherent risk measure:

$$\rho_t = (1-\lambda)\mathbb{E}[R_t] + \lambda \mathrm{CVaR}_\alpha(R_t) \tag{9}$$

The parameter $\lambda \in [0,1)$ can be tuned for a compromise between optimizing on average and considering risk aversion. This generalization still allows coming up with modified selling probabilities.

Theorem 2 *The risk-averse dynamic pricing problem $V_t^{\alpha,\lambda}(c)$ optimizing a combination of expected value and CVaR (9) can be transformed to a risk-neutral dynamic pricing problem $\tilde{V}_t^{\alpha,\lambda}(c)$ according* *to (3) with prices $\tilde{p}_t^{\alpha,\gamma} \in \mathcal{P}_t = [0, p_t^\infty]$ and a modified selling probability that distinguishes two cases depending on whether the event of selling is considered in CVaR's expectation:*

$$\tilde{d}_t^{\alpha,\gamma}(p) = \begin{cases} (1-\lambda) \cdot d_t(p) + \lambda\big(1 - \big(1 - d_t(p)\big)/\alpha\big), & d_t(p) \geq 1 - \alpha \\ (1-\lambda) \cdot d(p) & , \ 0 \leq d_t(p) < 1 - \alpha \end{cases} \quad \forall\, t \tag{10}$$

Then, the following holds:

a) *The objective values of both problems are equal: $V_t^{\alpha,\lambda}(c) = \tilde{V}_t^{\alpha,\lambda}(c)\ \forall\, t, c$.*

b) *The optimal risk-averse price $p_{t,c}^{\alpha,\lambda}$ in all states (t, c) is given by the risk-neutral problem with $\tilde{d}_t^{\alpha,\lambda}(p)$ (and vice versa): $p_{t,c}^{\alpha,\lambda} = \tilde{p}_t^{\alpha,\lambda} \quad \forall\, t, c$*

Proof: The proof of this theorem is given in Appendix A.4.

Lemma 4 *The transformation of the selling probability described in Theorem 2 preserves the first standard assumption, that is, $-\frac{\tilde{d}_t''(p)}{\tilde{d}_t'(p)} < -2\frac{\tilde{d}_t'(p)}{\tilde{d}_t(p)}\ \forall\, t \in \{1, 2, \dots, T\}, p \in (0, p_t^\infty)$ (see Section 3.1). Thus, the corresponding optimal price $p_{t,c}^{\alpha,\lambda}$ is well defined.*

Proof: It is necessary to show that the assumption is preserved for all p, i.e. that the revenue function is concave in demand. It is easy to see that it is preserved in the second case. The proof for the first case is based on the fact that the modified selling probability is a weighted sum of the original probability (for which the assumption holds) and the probability modified according to Theorem 1 (for which we know from Lemma 1 that the assumption holds). Thus, the resulting revenue function is the sum of two concave functions, which again is concave.

This shows how to transform a risk-averse dynamic pricing problem considering a combination of CVaR and expected value into a classical risk-neutral one. Although the modification of the selling probability now distinguishes two cases, it is important to recognize that the transformation is still static in the sense that it can be done offline in advance. The transformed selling probability is a (time-dependent) function of the original selling probability only. It does not depend on the state or the solution process. Thus, we still obtain a standard dynamic pricing problem, although with slightly more complicated modified selling probabilities, which might render the usage of existing algorithms more difficult. Nonetheless, all standard structural results including the monotonicities discussed in the previous subsection carry over.

4.4 Relation of risk-averse and risk-neutral solution

In this section, we again focus on CVaR only and compare a risk-averse and a risk-neutral decision maker who solve problem (7) and (3), respectively. Note that in both problems, the original selling probabilities $d_t(p)$ are used. The transformed problem described in Section 4.1 and the modified selling probabilities are not considered here. In particular, we show that the risk-averse and the risk-neutral solution are proportional under certain conditions.

Proposition 3 *If the distribution of the willingness-to-pay $F_t(p)$ and its density $f_t(p)$ satisfy the following conditions for a given risk level α, for all t, for all $p \in \mathcal{P}_t$, and $\exists y \in \mathbb{R}$ such that*

a) $F_t(\alpha^y p) = \alpha \cdot F_t(p)$ *and*

b) $f_t(\alpha^y p) = \alpha^{1-y} \cdot f_t(p)$,

then it holds that $p_{t,c}^{\alpha} = \alpha^y \cdot p_{t,c}^1$ and $V_t^{\alpha}(c) = \alpha^y \cdot V_t^1(c)$ for all c, t.

Proof: The proof of this proposition is given in Appendix A.4.

Remark 1 Proposition 3 implies that the optimal price $p_{t,c}^{\alpha}$ increases in the probability level α.

Proposition 3 shows a surprisingly straightforward connection between the optimal policies and objective values. The optimal price $p_{t,c}^{\alpha}$ a risk-averse firm sets is linear in the price $p_{t,c}^1$ a risk-neutral firm sets with a proportionality factor of α^y, where α is the given risk level and y follows from the structure of the distribution of the willingness-to-pay $F_t(p)$. The same relation holds for the objective values. This is especially remarkable in the context of literature that suggests risk-averse decision makers to heuristically consider risk by discounting. For example, Huang and Chang (2011) and Koenig and Meissner (2015b) propose capacity control approaches based on discounting opportunity cost. Different to this literature, we provide the optimal discount factor, i.e., α^y together with a sound theoretical foundation.

Several distributions satisfy the conditions of Proposition 3. One example is the uniform distribution. Then, we have $y = 1$ and obtain the linear transformation $p_{t,c}^{\alpha} = \alpha \cdot p_{t,c}^1$ and $V_t^{\alpha}(c) = \alpha \cdot V_t^1(\mathrm{c})$. Thus, the optimal price $p_{t,c}^{\alpha}$ and the value of the objective function $V_t^{\alpha}(c)$ are calculated by simply discounting $p_{t,c}^1$ and $V_t^1(c)$ with the discount factor (aka risk level) α, respectively. Two more examples are the distribution functions $F(p) = p^2$ $(y = 1/2)$ and $F(p) = \sqrt{p}$ $(y = 2)$ with p normalized to $[0, 1]$. We obtain $p_{t,c}^{\alpha} = \sqrt{\alpha} \cdot p_{t,c}^1$ and $p_{t,c}^{\alpha} = \alpha^2 \cdot p_{t,c}^1$, respectively. Among others, every distribution function with the structure $F(p) = b_0 \cdot p^{b_1}$ ($b_0, b_1 > 0$) fulfills the conditions of Proposition 3 and the corresponding linear transformation is given by $p_{t,c}^{\alpha} = \alpha^{\frac{1}{b_1}} \cdot p_{t,c}^{1}$.

Remark 2 Robust dynamic pricing assumes that the WTP distribution is not exactly known. For example, Lim and Shantikumar (2007) use relative entropy, a distance measure for distributions, to describe uncertainty and maximize worst-case expected revenue. Theorem 1 and Proposition 3 can also be interpreted in the context of robust dynamic pricing with a maximin objective. The probability measure $\mathbb{Q}$ is not exactly known, but belongs to an uncertainty set around a (forecasted) measure $\mathbb{P}$. This connection can be seen by writing CVaR's dual representation (4) as $\inf_{\mathbb{Q}\in\mathcal{U}}\{\mathbb{E}_{\mathbb{Q}}[R]\}$ with $\mathcal{U} = \left\{\mathbb{Q} \in \mathfrak{P} \middle| \mathbb{Q} = Z\mathbb{P}, \mathbb{E}[Z] = 1,\ 0 \leq Z \leq \frac{1}{\alpha}\right\}$. Then, our objective can be interpreted as an adversarial nature choosing the worst measure $\mathbb{Q}$ from the uncertainty set $\mathcal{U}$.

4.5 Example: U[0,1] distributed reservation prices

In this section, we consider a time-homogeneous uniform distribution of the willingness-to-pay on the interval $[0, 1]$. Obviously, all properties described in Sections 4.1 and 4.2 apply. The optimization problem becomes

$$V_t^{\alpha}(c) = \max_{p\in[0,1]} \min_{z_{t-1,c-1}, z_{t-1,c}} (1-p) \cdot z_{t-1,c-1} \cdot \left(p + V_{t-1}^{\alpha}(c-1)\right) + p \cdot z_{t-1,c} \cdot V_{t-1}^{\alpha}(c)$$

subject to

$$1 = p \cdot z_{t-1,c} + (1-p) \cdot z_{t-1,c-1}$$

$$0 \leq z_{t-1,c} \leq \frac{1}{\alpha},\ \ 0 \leq z_{t-1,c-1} \leq \frac{1}{\alpha}$$

To solve this optimization problem, we can either use Theorem 1 or Proposition 3. Following Theorem 1, we have $\tilde{d}_t^{\alpha}(p) = 1 - \frac{p}{\alpha}$ and $\mathcal{P}_t = [0, \alpha]$. Thus, we solve

$$\tilde{V}_t^{\alpha}(c) = \max_{p\in[0,\alpha]} \left(1 - \frac{p}{\alpha}\right) \cdot \left(p + \tilde{V}_{t-1}^{\alpha}(c-1)\right) + \frac{p}{\alpha} \cdot \tilde{V}_{t-1}^{\alpha}(c).$$

The optimal solution must obey the necessary and sufficient f.o.c., i.e., $1 - \frac{p}{\alpha} - \frac{1}{\alpha} \cdot \left(p + V_{t-1}^{\alpha}(c-1)\right) + \frac{1}{\alpha} \cdot V_{t-1}^{\alpha}(c) = 0$. After rearranging, the formula can be written as $p_{t,c}^{\alpha} = \frac{1}{2} \cdot \left(\alpha + V_{t-1}^{\alpha}(c) - V_{t-1}^{\alpha}(c-1)\right)$.

Following Proposition 3, we need to calculate $p_{t,c}^{\alpha} = \alpha \cdot p_{t,c}^{1}$ and $V_t^{\alpha}(c) = \alpha \cdot V_t^{1}(c)$. Thus, we first solve $V_t^1(c) = \max_{p\in[0,1]} (1-p) \cdot \left(p + V_{t-1}^1(c-1)\right) + p \cdot V_{t-1}^1(c)$. Again, the optimal solutions must satisfy the f.o.c.: $1 - p - \left(p + V_{t-1}^1(c-1)\right) + V_{t-1}^1(c) = 0$. After rearranging, the formula can be written as $p_{t,c}^1 = \frac{1}{2} \cdot \left(1 + V_{t-1}^1(c) - V_{t-1}^1(c-1)\right)$.

Now using the results of Proposition 3, we obtain that $p_{t,c}^{\alpha} = \alpha \cdot p_{t,c}^1 = \frac{1}{2} \cdot \left(\alpha + \alpha \cdot V_{t-1}^1(c) - \alpha \cdot V_{t-1}^1(c-1)\right) = \frac{1}{2} \cdot \left(\alpha + V_{t-1}^{\alpha}(c) - V_{t-1}^{\alpha}(c-1)\right)$ and both approaches provide the same optimal solution.

To summarize, we state: The solution of the risk-averse problem (7) with $U[0,1]$ distributed WTP is given by $p_{t,c}^{\alpha} = \frac{1}{2} \cdot \left(\alpha + \Delta_{t-1,c}^{\alpha}\right)$. The corresponding optimal objective value is $V_t^{\alpha}(c) = V_{t-1}^{\alpha}(c-1) + \frac{1}{\alpha} \cdot \left(p_{t,c}^{\alpha}\right)^2$.

5 Numerical studies

In this section, we consider a company (think of an airline) that sells a fixed capacity during a given selling horizon. We numerically compare the results of our approach to two common risk-averse approaches from literature. We compare the mechanisms by evaluating the results of the policies in terms of expected value and standard deviation. By varying their parameters, we are able to analyze the tradeoff between risk aversion and maximizing expected revenue. In Subsection 5.1, we stick to the standard setting of dynamic pricing and allow selling C units at most. We then extend the studies with a change in the setting, now allowing overbooking, as is common for example in the airline industry (Subsection 5.2). In particular, we investigate:

- *CVaR* is the approach we proposed in model (7).
- *DiscOC* is a heuristic mechanism derived from the risk-neutral dynamic program by taking opportunity costs only partially into account when determining the price to set. This mechanism was developed by Huang and Chang (2011) and investigated by Koenig and Meissner (2015a) in the context of capacity control and can easily be adapted to dynamic pricing. The risk-neutral dynamic program (3) is changed to $V_t^{\beta}(c) = d_t\left(p_{t,c}^{\beta}\right) \cdot \left(p_{t,c}^{\beta} + V_{t-1}^{\beta}(c-1)\right) + \left(1 - d_t\left(p_{t,c}^{\beta}\right)\right) \cdot V_{t-1}^{\beta}(c)$ with $p_{t,c}^{\beta} =$

$\underset{p\in\mathcal{P}_t}{\operatorname{argmax}}\, d_t(p)\cdot\left(p-\beta\cdot\left(V_t^{\beta}(c)-V_t^{\beta}(c-1)\right)\right)$ and the discounting factor $\beta\in[0,1]$. Prices are lower than in a risk-neutral setting, i.e. for $\beta=1$, and thus leading to a higher selling probability.

- *ExpUt* is a mechanism that maximizes the utility of total revenue R given by the exponential utility function $-e^{-\gamma R}$ with $\gamma\in[0,\infty)$. Among others, Barz and Waldmann (2007) and Li and Zhuang (2009) used this popular approach to risk aversion (see also Gönsch 2017 for more references). The dynamic program is given by $V_t^{\gamma}(c)=\max_{p\in\mathcal{P}_t}\left(1-d_t(p)\right)\cdot V_{t-1}^{\gamma}(c)+d_t(p)\cdot\left(e^{-\gamma p}\cdot V_{t-1}^{\gamma}(c-1)\right)$ with boundary conditions $V_0^{\gamma}(c)=-1$ for $c\geq 0$ and $V_t^{\gamma}(c)=-\infty$ for $c<0$. It can be shown that for $\gamma\to 0$ this approach approximates a maximization of $\mathbb{E}[R]-\frac{\gamma}{2}Var(R)$, and, thus a risk-neutral objective in the limit; for $\gamma\to\infty$ it reduces to a worst-case optimization (see Barz and Waldmann 2007).

Every mechanism produces a policy containing (optimal) selling prices for every possible state. We generated 10,000 customer streams in advance and applied the policies obtained from the mechanisms to the same streams. A simulation run of one customer stream mirrors a whole sales process with a price stated in each period according to the policy of the investigated mechanism and observing the arriving costumer's decision before moving on to the next period. Finally, the 10,000 outcomes for one policy provide the basis for estimating its expected revenue and standard deviation.

Although not every evaluation is shown in the following, we analyzed $T=10$ and $C=1,\dots,10$ using U[0,1] distributed WTP. We generated several policies with every mechanism by varying the parameter that describes the degree of risk aversion. Particularly, we choose $\alpha=0.01,0.02,\dots,1$ for *CVaR*, the same values as discount factor β for *DiscOC* and $\gamma=0.05,0.1,\dots,1,1.25,\dots,20$ for *ExpUt*.

5.1 Tradeoff between risk and revenue without overbooking

As mentioned, every evaluation of a policy results in 10,000 revenues that can be used to estimate the expected value and the standard deviation resulting from the underlying policy. These values are shown in the following figure. By varying the risk parameter (α, β, and γ), we generate multiple policies for every mechanism. We evaluate the policies and connect the data points resulting from the same mechanism using consecutive parameter values to a curve in Figure 3.

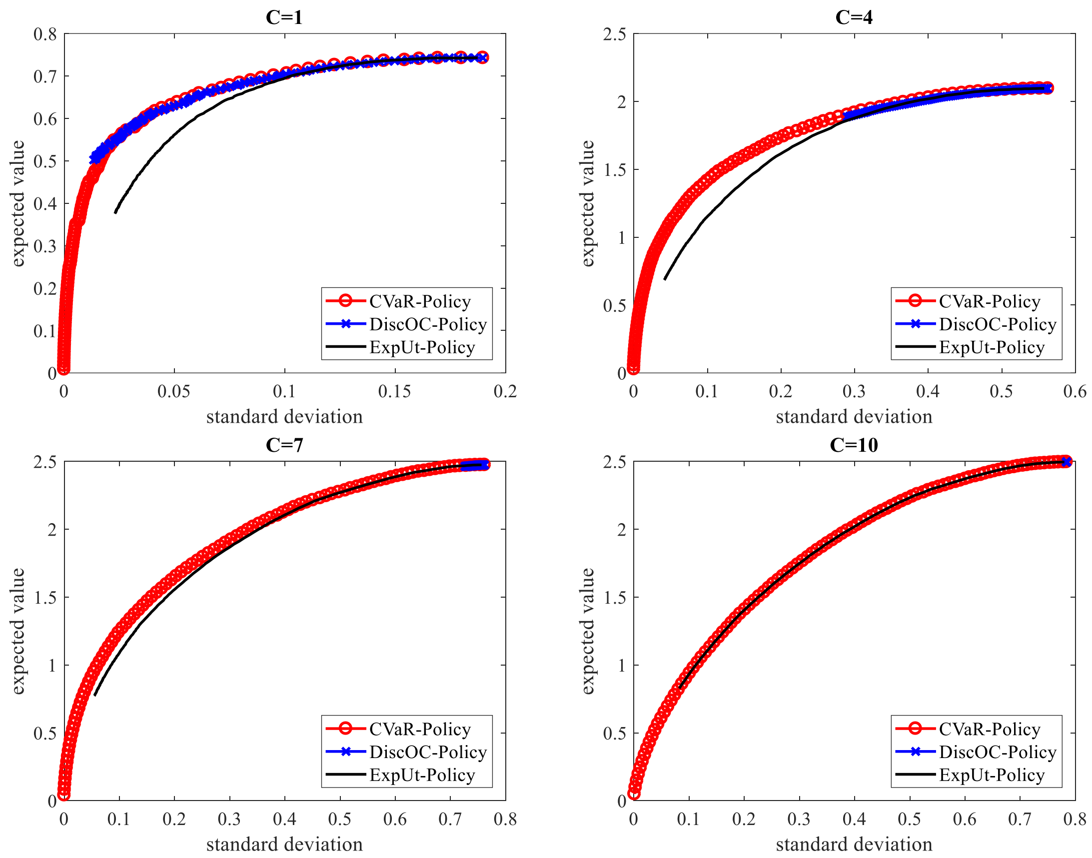

Figure 3: Tradeoff between expected value and standard deviation with $T = 10$ and $C = 1$ (upper left), $C = 4$ (upper right), $C = 7$ (lower left) and $C = 10$ (lower right)

For every mechanism, a higher degree of risk aversion, i.e. a lower α, a lower β, or a higher γ, leads to a lower standard deviation and a lower expected value (except *DiscOC* for $C = 10$). The amount of capacity has an impact on the performance of the mechanisms. While *CVaR* has an overall competitive performance, *ExpUt* is (weakly) dominated most of the time by *CVaR* (and partially by *DiscOC)*. *DiscOC* excels in the case of a high scarcity ($C = 1$). However, the usability of *DiscOC* vanishes around $C \geq 7$ as can be seen in Figure 3 where *DiscOC* only attains a very small interval of standard deviations for $C = 7$ and $C = 10$. This effect and the convergence of *ExpUt* to *CVaR* can be explained as follows: Opportunity costs are non-increasing in capacity and for $C = 7$ they are already nearly zero. Thus, varying the discount factor of the opportunity costs, i.e. β, impacts the policy resulting from *DiscOC* only slightly. Consequently, there are only small differences between the expected value and standard deviation resulting from *DiscOC* with different discount factors β. Simultaneously, the policies of *ExpUt* and *CVaR* are hardly affected by opportunity costs as that they are nearly zero. Thus, the prices are changing very little over the selling horizon. Contrary to *DiscOC*, the risk parameters of *CVaR* and *ExpUt*, i.e. α and γ respectively, still have a high impact on the corresponding nearly static policy. All in all, we can say that for every α there is a γ such that the resulting nearly static policies are varying only to a small degree, and thus, expected value and standard deviation are nearly the same. For $C = 10$, opportunity costs are zero and stay zero for the whole selling process as we always have at least one capacity per customer left. Thus, *DiscOC* is completely independent of β and always leads to the risk-neutral policy. Moreover, every mechanism now has a static policy and we can always find for every α a corresponding γ such that *CVaR* and *ExpUt* produce exactly the same policy resulting in the same tradeoffs between expected value and standard deviation.

To check the robustness of the aforementioned results, we next consider a bigger setting with a longer time horizon. In particular, Figure 4 shows the tradeoff between expected value and standard deviation with $T = 50$ for the two capacities $C = 20$ (left column) and $C = 35$ (right column). In addition to the uniform distribution used up to now (top row), we also consider the additional distributions $F(p) = \sqrt{p}$ (middle row) and $F(p) = p^2$ (bottom row). By and large, the aforementioned results are still valid. As the time horizon was scaled up with a factor of five, we briefly compare the same time/capacity ratios, that is, $C = 20$ with $C = 4$ and $C = 35$ with $C = 7$. Regarding the uniform distribution, almost nothing changed. For $C = 20$, *DiscOC* spans a slightly larger range of standard deviations and dominates *CVaR* at its lower end. For $F(p) = \sqrt{p}$, *DiscOC* is basically not applicable and *CVaR* dominates *ExpUt*. For $F(p) = p^2$ and $C = 20$, *DiscOC* performs very good, attains all standard deviations, and dominates all other methods for the lower third of the interval. For $C = 35$, it is again basically useless. *CVaR* still attains all standard deviations and strictly dominates *ExpUt* for about the lower half of the interval.

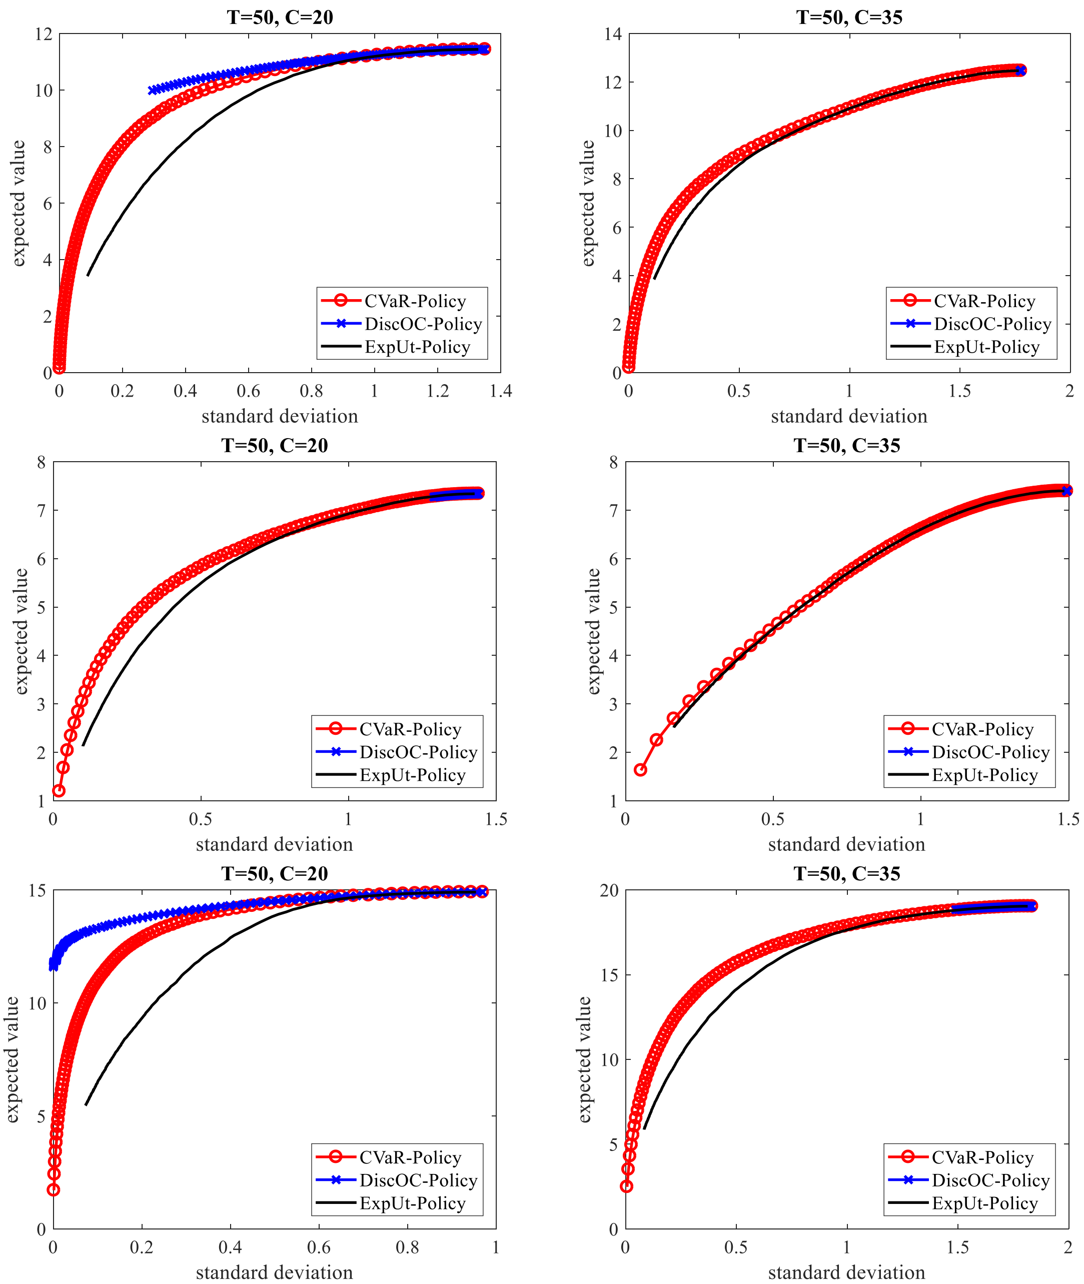

Figure 4: Tradeoff between expected value and standard deviation with $T = 50$, $C = 20$ (left column), and $C = 35$ (right column) for different distributions of the willingness to pay: $F(p) = p$ (uniform distribution, top row), $F(p) = \sqrt{p}$ (middle row), and $F(p) = p^2$ (bottom row)

Finally, we briefly discuss runtime. From a theoretical perspective, CVaR is very efficient. The number of states is the same as in risk-neutral dynamic pricing. Likewise, for a variety of demand distributions, there is a closed-form solution for the one-stage optimization problem solved in each state. As expected, Figure 5 shows that the algorithm is very fast and runtime increases only linearly in T and C, that is, we have $O(T \cdot C)$. The minor deviation of actual runtime from the predicted scaling is probably due to operating system tasks on the CPU and similar technical issues. We did not depict the runtime for *DiscOC* and *ExpUt*, as it is comparable. For a given β, *DiscOC* considers the same states and also enjoys a closed-form solution. *ExpUt* also considers the same states and, thus, has a comparable runtime. However, we solved the one-stage problem in *ExpUt* numerically, which scales runtimes by a constant factor.

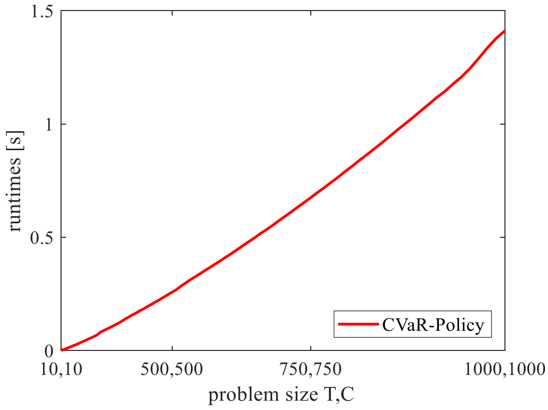

Figure 5: Runtime of *CVaR* depending on problem size

5.2 Tradeoff between risk and revenue with overbooking

We now add the possibility to overbook to our basic setting ($T = 10$). This means that the firm is allowed to sell more units of the product than capacity is available. It may be beneficial because we now assume that buyers have a given show rate determining the probability that they claim their product in the service period after the selling horizon. Whenever a customer's claim cannot be met, the firm has to pay a penalty. To optimize the selling horizon, the firm anticipates that some customers might not show up. Technically speaking, the boundary conditions of the model must be adapted to consider show rate and penalty costs if more capacity has been sold then available. Thus, additional states have to be taken into account. In particular, now states with $t \in \{1, 2, \dots, T\}$ and $c \in \{C - T, C - T - 1, \dots, 0, 1, 2, \dots, C\}$ are possible, and thus the model has to deal with negative capacities. The boundary conditions have to be changed to the effect that, for $t = 0$ and $c < 0$, the number of customers' claims depends on the show rate and the number of sold products and follows a binomial distribution. For every possible number of claims, the penalty costs for every claim exceeding capacity have to be considered. To evaluate these states, the mechanisms proceed as follows for the last period $t = 0$: *CVaR* calculates the conditional value at risk of the penalty cost's distribution, *DiscOC* uses the expected penalty and *ExpUt* the expected utility of the penalty costs. With the new boundary conditions and the state space adapted to overbooking, the mechanisms can use their underlying statewise optimization during the selling horizon ($t \in \{1, 2, \ldots, T\}$). Only for *DiscOC*, a small change is necessary to reflect risk aversion. We now increase the opportunity costs whenever there is no capacity left, i.e. we use $p_{t,c}^{\beta} = \underset{p \in \mathcal{P}_t}{\operatorname{argmax}}\, d_t(p) \cdot \left(p - \frac{1}{\beta} \cdot \left(V_t^{\beta}(c) - V_t^{\beta}(c-1)\right)\right)$ if $c \leq 0$. By doing so, compared to the risk-neutral approach, we increase the selling probability while $c > 0$ and decrease it as soon as we run out of capacity because less overbooking is intuitively risk-averse.

In Figure 6, we show the results of our simulation study using a show rate of 0.8 and penalty costs of 1, i.e. the highest possible WTP. We varied both parameters in further studies but found the results of no further interest, and thus omitted them in this paper.

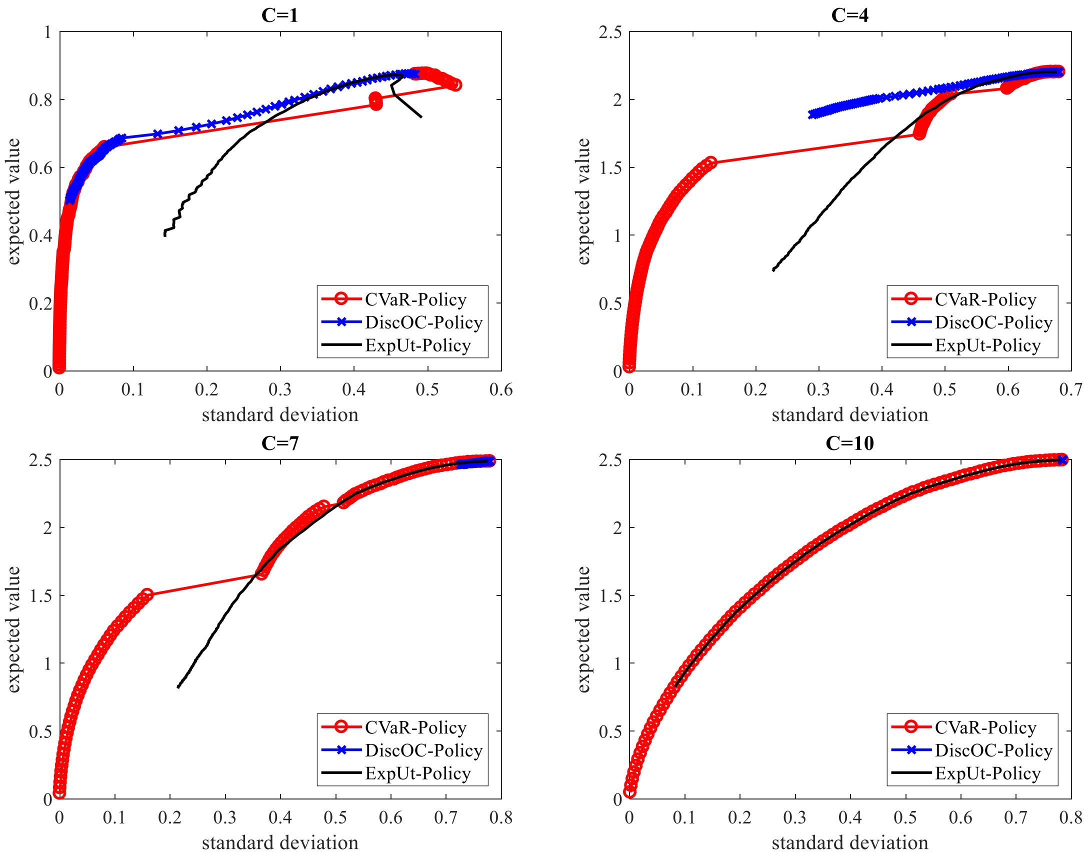

Figure 6: Tradeoff between expected value and standard deviation with overbooking with $T = 10$ and $C = 1$ (upper left), $C = 4$ (upper right), $C = 7$ (lower left) and $C = 10$ (lower right)

Obviously, *CVaR* has the greatest change in its structure. As we can see, there are several gaps in its curve compared to the previous subsection without overbooking. A closer look in the data shows that *CVaR* has for every α a certain amount of overbooking it permits. When this amount changes, this structural change leads to a gap in the curve. In particular, for $C = 1$ overbooking is completely banned for $\alpha \in [0.01, 0.8]$, partially allowed up to one unit for $\alpha \in [0.81, 0.84]$ and fully allowed, i.e. up to 9 units, for $\alpha \in [0.85, 1]$ ($C = 4$: no overbooking for $\alpha \in [0.01, 0.57]$, up to 1 unit for $\alpha \in [0.58, 0.76]$, up to 2 units for $\alpha \in [0.77, 0.84]$ and up to 6 units for $\alpha \in [0.85, 1]$; $C = 7$: 0 for $\alpha \in [0.01, 0.4]$, up to 1 unit for $\alpha \in [0.41, 0.63]$, up to 2 units for $\alpha \in [0.64, 0.76]$ and up to 3 units for $\alpha \in [0.77, 1]$; $C = 10$: no overbooking possible). As the conditional value at risk only considers the α-worst outcomes, the states with $t = 0$ and $c < 0$ are highly negative valued because the mechanism focuses on the cases where most if not all customers claim their overbooked products.

As an example, consider the following setting: $C = 1$, $\alpha = 0.8$, show rate 0.8, penalty costs 1. We start with calculating the boundary conditions for $t = 0$ and $c = 0$ as well as $c = -1$, i.e. $V_0^\alpha(0)$ and $V_0^\alpha(-1)$. Obviously, $V_0^\alpha(0) = 0$. To calculate $V_0^\alpha(-1)$, we have to consider the probabilities that both customers ($C - c = 1 - (-1) = 2$) claim their purchased product, that only one customer shows up to claim his product or no customer claims a product at all. The probabilities for these events are 0.64, 0.32 and 0.04, respectively. The α-worst outcomes are two and one claims with penalties of 1 and 0, respectively. Thus, we have $V_0^\alpha(-1) = CVar_{0.8} = \frac{0.64}{0.8} \cdot (-1) + \frac{0.16}{0.8} \cdot 0 = -0.8$. Now, when optimizing $V_1^\alpha(0)$, we consider opportunity costs of 0.8. As this is greater than or equal to $\alpha = 0.8$, the optimal price is given by p_t^∞. After similar considerations, we find that p_t^∞ is the optimal solution for every $V_t^\alpha(0)$ with $t = 1, \dots, 10$, and thus overbooking is completely avoided in this example.

The difference between no overbooking and up to 1 unit is noticeable for *CVaR*. While the expected value slightly increases, there is a huge increase in the standard deviation (first gap from left). By comparison, *CVaR* tends to set lower prices than the other mechanisms, especially after the first gap. Thereby, in the simulation, *CVaR* sells more and has to pay more often than *DiscOC* and *ExpUt* the penalty costs. Moreover, the impact of the penalty costs is higher for *CVaR* as the prices earned for every overbooked capacity are lower. Together, this leads to a significantly higher increase in standard deviation. Another interesting observation is that for $C = 1$ and $\alpha \geq 0.85$ the expected value is increasing and the standard deviation is decreasing in α. This seems contra intuitive as a lower risk aversion, i.e. a higher α, is expected to generate a riskier policy resulting in a higher standard deviation. The explanation of this observation is as follows: A higher α leads to higher prices. Although higher prices lead to an increased variation of the revenues, the overall standard deviation decreases as the selling probability, and thus the probability of paying penalty costs, decreases.

Comparing the performance of the mechanisms, we can observe that *CVaR* is a good choice whenever capacity is quite high ($C = 7$ and $C = 10$) or risk aversion is strong (left part of the curves in Figure 6, upper left and upper right). *DiscOC* performs very well whenever capacity is scarce, while *ExpUt* is mostly dominated by one or two of the other mechanisms.

6 Conclusion and managerial implications

In this paper, we have shown how the risk-averse dynamic pricing problem with nested CVaR can be transformed to an equivalent, standard dynamic pricing problem maximizing expected revenue. The transformation is easy to apply and, thus, the risk-averse problem is easy to compute. Besides time consistency, this is a major advantage compared to other approaches to include risk in dynamic pricing, which usually considerably increase the computational burden. As a consequence of the transformation, the well-known monotonicities carry over and new, risk-related monotonicities have been shown. Finally, under certain conditions regarding demand, the risk-neutral and risk-averse solutions are proportional regarding objective value and optimal prices.

In a numerical study, we have analyzed the tradeoff between expected revenue and standard deviation for dynamic pricing with and without overbooking. In this regard, the new approach performs very well compared to two standard approaches from literature.

Our results have several direct implications for practice. First and foremost, they allow to use existing, standard dynamic pricing algorithms and systems for a straightforward, theoretically sound, risk-averse dynamic pricing. Second, the modification of the selling probabilities is in line with intuitive behavior of many risk-averse decision makers who often use “conservative” estimates of selling probabilities. Finally, the proportionality of the risk-averse and the risk-neutral solution not only provides an easy alternative to solve the risk-averse dynamic pricing problem. It also shows the objective risk-averse managers, who ‘heuristically’ discount optimal prices, may implicitly optimize.

References

Ahmed, S., U. Çakmak, A. Shapiro. 2007. Coherent risk measures in inventory problems. *European Journal of Operational Research*, 182 (1), 226–238.

Artzner, P., F. Delbaen, J.-M. Eber, D. Heath. 1999. Coherent measures of risk. *Mathematical Finance*, 9 (3), 203–228.

Artzner, P., F. Delbaen, J.-M. Eber, D. Heath, H. Ku. 2007. Coherent multiperiod risk adjusted values and Bellman's principle. *Annals of Operations Research*, 152 (1), 5–22.

Bamberg, G., M. Krapp. 2016. Is time consistency compatible with risk aversion? *Review of Managerial Science*, 10 (2), 195–211.

Barz, C., K. H. Waldmann. 2007. Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research*, 65 (3), 565–579.

Bellman, R. 1957. *Dynamic Programming*, 3rd ed. Princeton Univ. Press, Princeton.

Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing & Service Operations Management*, 5 (3), 203–229.

Bitran, G. R., S. V. Mondschein. 1997. Periodic pricing of seasonal products in retailing. *Management Science*, 43 (1), 64–79.

Cachon, G. P., M. A. Lariviere. 2001. Contracting to assure supply: how to share demand forecasts in a supply chain. *Management Science*, 47 (5), 629–646.

Caliskan-Demirag, O., Y. Chen, J. Li. 2011. Customer and retailer rebates under risk aversion. *International Journal of Production Economics*, 133 (2), 736–750.

Chen, Y., M. Xu, Z. G. Zhang. 2009. Technical note-a risk-averse newsvendor model under the cvar criterion. *Operations Research*, 57 (4), 1040–1044.

Cheng, L., Z. Wan, G. Wang. 2009. Bilevel newsvendor models considering retailer with CVaR objective. *Computers & Industrial Engineering*, 57 (1), 310–318.

Cheridito, P., F. Delbaen, M. Kupper. 2006. Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability*, 11 (3), 57–106.

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

Collado, R. A., D. Papp, A. Ruszczyński. 2012. Scenario decomposition of risk-averse multistage stochastic programming problems. *Annals of Operations Research*, 200 (1), 147–170.

den Boer, A. V. 2015. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, 20 (1), 1–18.

Feichtinger, G., R. Hartl. 1985. Optimal pricing and production in an inventory model. *European Journal of Operational Research*, 19 (1), 45–56.

Feng, Y., B. Xiao. 1999. Maximizing revenues of perishable assets with a risk factor. *Operations Research*, 47 (2), 337–341.

Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40 (8), 999–1020.

Gönsch, J. 2017. A survey on risk-averse and robust revenue management. *European Journal of Operational Research*, 263 (2), 337–348.

Gönsch, J., M. Hassler. 2014. Optimizing the conditional value-at-risk in revenue management. *Review of Managerial Science*, 8 (4), 495–521.

Gönsch, J., M. Hassler, R. Schur. 2018. Optimizing conditional value-at-risk in dynamic pricing. *OR Spectrum*, 40 (3), 711–750.

Gönsch, J., R. Klein, M. Neugebauer, C. Steinhardt. 2013. Dynamic pricing with strategic customers. *Journal of Business Economics*, 83 (5), 505–549.

Gotoh, J.-Y., Y. Takano. 2007. Newsvendor solutions via conditional value-at-risk minimization. *European Journal of Operational Research*, 179 (1), 80–96.

Hempenius, A. L. 1970. *Monopoly with random demand.* Rotterdam University Press, Rotterdam.

Hsieh, C.-C., Y.-T. Lu. 2010. Manufacturer's return policy in a two-stage supply chain with two risk-averse retailers and random demand. *European Journal of Operational Research*, 207 (1), 514–523.

Huang, K., K.-C. Chang. 2011. A model for airline seat control considering revenue uncertainty and risk. *Journal of Revenue and Pricing Management*, 10 (2), 161–171.

Koch, S., J. Gönsch, M. Hassler, R. Klein. 2016. Practical decision rules for risk-averse revenue management using simulation-based optimization. *Journal of Revenue and Pricing Management*, 15 (6), 468–487.

Koenig, M., J. Meissner. 2010. List pricing versus dynamic pricing: impact on the revenue risk. *European Journal of Operational Research*, 204 (3), 505–512.

Koenig, M., J. Meissner. 2015a. Risk management policies for dynamic capacity control. *Computers & Operations Research*, 59, 104–118.

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

Koenig, M., J. Meissner. 2016. Risk minimising strategies for revenue management problems with target values. *Journal of the Operational Research Society*, 67 (3), 402–411.

Kovacevic, R., G. Pflug. 2009. Time consistency and information monotonicity of multiperiod acceptability functionals. H. Albrecher, ed. *Advanced financial modelling*. de Gruyter, Berlin, 347–369.

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

Landsberger, M., I. Meilijson. 1985. Intertemporal price discrimination and sales strategy under incomplete information. *The RAND Journal of Economics*, 16 (3), 424–430.

Lariviere, M. A., E. L. Porteus. 2001. Selling to the newsvendor: an analysis of price-only contracts. *Manufacturing & Service Operations Management*, 3 (4), 293–305.

Levin, Y., J. McGill, M. Nediak. 2008. Risk in revenue management and dynamic pricing. *Operations Research*, 56 (2), 326–343.

Li, L. 1988. A Stochastic Theory of the Firm. *Mathematics of Operations Research*, 13 (3), 447–466.

Li, M. Z., W. Zhuang. 2009. Risk-sensitive dynamic pricing for a single perishable product. *Operations Research Letters*, 37 (5), 327–332.

Lim, A. E. B., J. G. Shanthikumar. 2007. Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research*, 55 (2), 198–214.

Luciano, E., L. Peccati, D. M. Cifarelli. 2003. VaR as a risk measure for multiperiod static inventory models. *International Journal of Production Economics*, 81–82, 375–384.

Maceira, M., L. Marzano, D. Penna, A. L. Diniz, T. C. Justino. 2015. Application of CVaR risk aversion approach in the expansion and operation planning and for setting the spot price in the Brazilian hydrothermal interconnected system. *International Journal of Electrical Power & Energy Systems*, 72, 126–135.

Paschalidis, I., J. N. Tsitsiklis. 2000. Congestion-dependent pricing of network services. *IEEE/ACM Transactions on Networking*, 8 (2), 171–184.

Pflug, G. 2000. Some remarks on the value-at-risk and the conditional value-at-risk. S. Uryasev, ed. *Probabilistic constrained optimization*. Springer, 272–281.

Pflug, G., A. Pichler. 2014. *Multistage Stochastic Optimization.* Springer, Cham.

Pflug, G., A. Pichler. 2016. Time-inconsistent multistage stochastic programs: martingale bounds. *European Journal of Operational Research*, 249 (1), 155–163.

Phillips, R. 2005. *Pricing and Revenue Optimization.* Stanford University Press, Stanford.

Philpott, A., V. de Matos. 2012. Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research*, 218 (2), 470–483.

Philpott, A., V. de Matos, E. Finardi. 2013. On solving multistage stochastic programs with coherent risk measures. *Operations Research*, 61 (4), 957–970.

Pousinho, H., V. Mendes, J. Catalão. 2012. Scheduling of a hydro producer considering head-dependency, price scenarios and risk-aversion. *Energy Conversion and Management*, 56, 96–103.

Riedel, F. 2004. Dynamic coherent risk measures. *Stochastic Processes and their Applications*, 112 (2), 185–200.

Rockafellar, R. T., S. Uryasev. 2000. Optimization of conditional value-at-risk. *Journal of Risk*, 2, 21–42.

Rudloff, B., A. Street, D. M. Valladão. 2014. Time consistency and risk averse dynamic decision models: definition, interpretation and practical consequences. *European Journal of Operational Research*, 234 (3), 743–750.

Ruszczyński, A. 2010. Risk-averse dynamic programming for markov decision processes. *Mathematical Programming*, 125 (2), 235–261.

Ruszczyński, A., A. Shapiro. 2006. Conditional risk mappings. *Mathematics of Operations Research*, 31 (3), 544–561.

Schlosser, R. 2015. A stochastic dynamic pricing and advertising model under risk aversion. *Journal of Revenue & Pricing Management*, 14 (6), 451–468.

Shapiro, A. 2009. On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters*, 37 (3), 143–147.

Shapiro, A. 2011. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209 (1), 63–72.

Shapiro, A., D. Dentcheva, A. P. Ruszczyński. 2009. *Lectures on stochastic programming*: *Modeling and theory*. Society for Industrial and Applied Mathematics, Philadelphia.

Shapiro, A., W. Tekaya, J. P. da Costa, M. P. Soares. 2013. Risk neutral and risk averse stochastic dual dynamic programming method. *European Journal of Operational Research*, 224 (2), 375–391.

Shapiro, A., K. Ugurlu. 2016. Decomposability and time consistency of risk averse multistage programs. *Operations Research Letters*, 44 (5), 663–665.

Stokey, N. L. 1979. Intertemporal price discrimination. *The Quarterly Journal of Economics*, 93 (3), 355–371.

Street, A. 2010. On the Conditional Value-at-Risk probability-dependent utility function. *Theory and Decision*, 68 (1-2), 49–68.

Talluri, K. T., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management.* Springer, New York.

von Neumann, J., O. Morgenstern. 1944. *Theory of Games and Economic Behavior.* Princeton University Press, Princeton.

Weber, S. 2006. Distribution-invariant risk measures, information, and dynamic consistency. *Mathematical Finance*, 16 (2), 419–441.

Wilson, C. A. 1988. On the optimal pricing policy of a monopolist. *The Journal of Political Economy*, 164–176.

Xue, W., L. Ma, H. Shen. 2015. Optimal inventory and hedging decisions with CVaR consideration. *International Journal of Production Economics*, 162, 70–82.

Ziya, S., H. Ayhan, R. D. Foley. 2004. Relationships among three assumptions in revenue management. *Operations Research*, 52 (5), 804–809.

Appendix: Time-Consistent, Risk-Averse Dynamic Pricing

A.1 Proof of Lemma 1 and Theorem 1

First, we show that the standard assumptions regarding the selling probabilities are preserved (Lemma 1). Second, we show statements a) and b) by induction (Theorem 1).

A.1.1 Preservation of standard assumptions regarding selling probabilities (Lemma 1)

First, we show the preservation of all standard assumptions regarding the selling probabilities. They have in common that the probability function $d_t(p)$ is two times continuously differentiable and strictly monotonically decreasing on $(0, p_t^\infty)\ \forall\, t \in \{1, 2, \dots, T\}$. Moreover, at least one of the following properties holds:

1. $-\frac{d_t''(p)}{d_t'(p)} < -2\frac{d_t'(p)}{d_t(p)} \qquad \forall\, t \in \{1, 2, \dots, T\}, p \in (0, p_t^\infty)$
2. $-\frac{d_t''(p)}{d_t'(p)} < -\frac{d_t'(p)}{d_t(p)} + \frac{1}{p} \quad \forall\, t \in \{1, 2, \dots, T\}, p \in (0, p_t^\infty)$

Obviously, it holds that $\tilde{d}_t^\alpha(p) = 1 - \big(1 - d_t(p)\big)/\alpha$ is also two times continuously differentiable and strictly monotonically decreasing on $(0, p_t^\infty)$. On this interval it follows

1. $-\frac{(\tilde{d}_t^\alpha)''(p)}{(\tilde{d}_t^\alpha)'(p)} = -\frac{d_t''(p)}{d_t'(p)} < -2\frac{d_t'(p)}{d_t(p)} = -2\frac{\alpha\cdot(\tilde{d}_t^\alpha)'(p)}{d_t(p)} \le -2\frac{(\tilde{d}_t^\alpha)'(p)}{d_t(p)} \le -2\frac{(\tilde{d}_t^\alpha)'(p)}{\tilde{d}_t^\alpha(p)}$
2. $-\frac{(\tilde{d}_t^\alpha)''(p)}{(\tilde{d}_t^\alpha)'(p)} = -\frac{d_t''(p)}{d_t'(p)} < -\frac{d_t'(p)}{d_t(p)} + \frac{1}{p} \le -\frac{(\tilde{d}_t^\alpha)'(p)}{\tilde{d}_t^\alpha(p)} + \frac{1}{p}$

As $\tilde{d}_t^\alpha(p)$ is monotone decreasing and thus, $\big(\tilde{d}_t^\alpha\big)'(p) < 0$, the last inequality of 1 holds if and only if $\tilde{d}_t^\alpha(p) \le d_t(p)$. By definition of $\tilde{d}_t^\alpha(p)$, this is true as $\tilde{d}_t^\alpha(p) \le d_t(p) \Leftrightarrow 1 - \big(1 - d_t(p)\big)/\alpha \le d_t(p) \Leftrightarrow \alpha - \big(1 - d_t(p)\big) \le \alpha \cdot d_t(p) \Leftrightarrow \alpha \cdot \big(1 - d_t(p)\big) \le 1 - d_t(p) \Leftrightarrow \alpha \le 1$. The last inequality of 2. holds with the same arguments.

A.1.2 Proof of a) and b) by induction (Theorem 1)

Next, we show by induction that the optimal objective values of both optimization problems are equal and, simultaneously, that the optimal risk-averse price $p_{t,c}^{\alpha}$ is given by the (transformed) risk neutral problem (and vice versa). For $t = 0$, the boundary conditions of both optimization problems are the same and, thus, the objective values are the same (induction basis). In the induction step, we assume that $V_{t-1}^{\alpha}(c) = \tilde{V}_{t-1}^{\alpha}(c)\ \forall\ c$ (induction hypothesis, IH) and show that this equality also holds for t. Simultaneously, we confirm that $p_{t,c}^{\alpha} = \begin{cases} \tilde{p}_{t,c}^{\alpha},\ \tilde{p}_{t,c}^{\alpha} \in \left[0, p_t^{\alpha,\infty}\right) \\ p_t^{\infty},\ \tilde{p}_{t,c}^{\alpha} = p_t^{\alpha,\infty} \end{cases}\quad \forall\ t, c.$

To be precise, we distinguish two cases regarding the optimal risk-averse price $p_{t,c}^{\alpha} \in [0, p_t^{\infty}]$: $p_{t,c}^{\alpha} \in \left[0, p_t^{\alpha,\infty}\right)$ and $p_{t,c}^{\alpha} \in \left[p_t^{\alpha,\infty}, p_t^{\infty}\right]$. We examine both intervals and show the equality of the value functions in both cases, i.e. $V_t^{\alpha}(c) = \tilde{V}_t^{\alpha}(c)$. While this can be straightforwardly shown for $p_{t,c}^{\alpha} \in \left[p_t^{\alpha,\infty}, p_t^{\infty}\right]$, we have to take an indirect way to show this equality for $p_{t,c}^{\alpha} \in \left[0, p_t^{\alpha,\infty}\right)$. There, we start with showing that the optimal value of the risk-averse problem does not exceed the optimal value of the transformed, risk-neutral problem. Next, we show that this relation also holds vice versa. Finally, we confirm the relation of both optimal prices.

A.1.2.1 Proof: $V_t^{\alpha}(c) = \tilde{V}_t^{\alpha}(c)$ for $p_{t,c}^{\alpha} \in \left[p_t^{\alpha,\infty}, p_t^{\infty}\right]$

In this case, we have that $p_{t,c}^{\alpha} \in \left[p_t^{\alpha,\infty}, p_t^{\infty}\right]$, and thus, from the definition of $p_t^{\alpha,\infty}$, we have $d_t\left(p_{t,c}^{\alpha}\right) \leq 1 - \alpha$. This case occurs if and only if $\Delta_{t,c}^{\alpha} \geq p_t^{\alpha,\infty}$ with $\Delta_{t,c}^{\alpha} = V_{t-1}^{\alpha}(c) - V_{t-1}^{\alpha}(c-1) \underbrace{=}_{\mathrm{IH}} \tilde{V}_{t-1}^{\alpha}(c) - \tilde{V}_{t-1}^{\alpha}(c-1) = \tilde{\Delta}_{t,c}^{\alpha}$. Otherwise, we could find a price $p \in \left(\Delta_{t,c}^{\alpha}, p_t^{\alpha,\infty}\right)$ with a higher objective value (greater than $V_{t-1}^{\alpha}(c)$) and that would contradict our assumption. Back to our case, we have $V_t^{\alpha}(c) = V_{t-1}^{\alpha}(c)$ from the minimization calculating CVaR, and either the optimal risk-averse price is equal to p_t^{∞} (and thus a unique solution) or the whole interval $\left[\Delta_{t,c}^{\alpha}, p_t^{\infty}\right]$ is optimal. This distinction depends on whether the interval is empty or not, whereby an empty interval can only occur if the distribution of the willingness-to-pay varies over t. Note that because $p_{t,c}^{\alpha} = p_t^{\infty}$ is always an optimal solution we choose it and do not need to distinguish any longer. With the induction hypothesis, it holds that $\tilde{\Delta}_{t,c}^{\alpha} =$

$\Delta_{t,c}^{\alpha} \geq p_t^{\alpha,\infty}$ and hence, regarding the transformed problem, $\tilde{d}_t^{\alpha}\left(\tilde{\Delta}_{t,c}^{\alpha}\right) \leq 0$. Thus, with a similar reasoning as above, we have $\tilde{p}_{t,c}^{\alpha} = p_t^{\alpha,\infty}$ with $\tilde{d}_t^{\alpha}(p_t^{\alpha,\infty}) = 1 - \left(1 - d_{\mathrm{t}}(p_t^{\alpha,\infty})\right)/\alpha = 1 - (1 - 1 + \alpha)/\alpha = 0$. Consequently, it holds that

$$V_t^{\alpha}(c) = \min_{\substack{z_{t-1,c-1},\\ z_{t-1,c}}} \left(1 - d_t\left(p_{t,c}^{\alpha}\right)\right) \cdot z_{t-1,c} \cdot V_{t-1}^{\alpha}(c) + d_t\left(p_{t,c}^{\alpha}\right) \cdot z_{t-1,c-1} \cdot \left(p_{t,c}^{\alpha} + V_{t-1}^{\alpha}(c-1)\right)$$
$$= V_{t-1}^{\alpha}(c) \underbrace{=}_{\mathrm{IH}} \tilde{V}_{t-1}^{\alpha}(c) = \tilde{d}_t^{\alpha}\left(\tilde{p}_{t,c}^{\alpha}\right) \cdot \left(\tilde{p}_{t,c}^{\alpha} + \tilde{V}_{t-1}^{\alpha}(c-1)\right) + \left(1 - \tilde{d}_t^{\alpha}\left(\tilde{p}_{t,c}^{\alpha}\right)\right) \cdot \tilde{V}_{t-1}^{\alpha}(c) = \tilde{V}_t^{\alpha}(c).$$

Above, the first equality is the definition of the risk-averse value function without the maximum as $p_{t,c}^{\alpha}$ is the optimal price. The second follows from $z_{t-1,c-1} = 0$ and $z_{t-1,c} = 1/\left(1 - d_t\left(p_{t,c}^{\alpha}\right)\right) \leq 1/\alpha$, which is the solution of the minimization because $V_{t-1}^{\alpha}(c) \leq p_{t,c}^{\alpha} + V_{t-1}^{\alpha}(c-1)$. The third is the induction hypothesis. The fourth equality holds because we have $\tilde{p}_{t,c}^{\alpha} = p_t^{\alpha,\infty}$ from $\tilde{\Delta}_{t,c}^{\alpha} = \Delta_{t,c}^{\alpha} \geq p_t^{\alpha,\infty}$ and $\tilde{d}_t^{\alpha}(p_t^{\alpha,\infty}) = 0$ (see above). Finally, the fifth equality is the transformed value function without the maximization, as $\tilde{p}_{t,c}^{\alpha}$ is its optimal price.

A.1.2.2 Proof: $V_t^{\alpha}(c) \leq \tilde{V}_t^{\alpha}(c)$ for $p_{t,c}^{\alpha} \in \left[0, p_t^{\alpha,\infty}\right)$

If $p_{t,c}^{\alpha} \in \left[0, p_t^{\alpha,\infty}\right)$ holds, we have (remember that $p_{t,c}^{\alpha}$ deonotes the optimal price)

$$\begin{aligned} V_t^{\alpha}(c) = \min_{\substack{z_{t-1,c-1},\\ z_{t-1,c}}} & \left(1 - d_t\left(p_{t,c}^{\alpha}\right)\right) \cdot z_{t-1,c} \cdot V_{t-1}^{\alpha}(c) + d_t\left(p_{t,c}^{\alpha}\right) \cdot z_{t-1,c-1} \cdot \left(p_{t,c}^{\alpha} + V_{t-1}^{\alpha}(c-1)\right) \\ & \leq \frac{1 - d_t\left(p_{t,c}^{\alpha}\right)}{\alpha} \cdot V_{t-1}^{\alpha}(c) + \left(1 - \frac{1 - d_t\left(p_{t,c}^{\alpha}\right)}{\alpha}\right) \\ & \cdot \left(p_{t,c}^{\alpha} + V_{t-1}^{\alpha}(c-1)\right) \underbrace{=}_{IH} \frac{1 - d_t\left(p_{t,c}^{\alpha}\right)}{\alpha} \cdot \tilde{V}_{t-1}^{\alpha}(c) + \left(1 - \frac{1 - d_t\left(p_{t,c}^{\alpha}\right)}{\alpha}\right) \\ & \cdot \left(p_{t,c}^{\alpha} + \tilde{V}_{t-1}^{\alpha}(c-1)\right) \leq \tilde{V}_t^{\alpha}(c) \end{aligned}$$

where the first inequality follows with the following feasible solution for the minimization problem: $z_{t-1,c} = \frac{1}{\alpha}$ and $z_{t-1,c-1} = \left(1 - \frac{1 - d_t\left(p_{t,c}^{\alpha}\right)}{\alpha}\right)/d_t\left(p_{t,c}^{\alpha}\right)$. As $p_{t,c}^{\alpha} \in \left[0, p_t^{\alpha,\infty}\right)$, the optimal risk-averse price is a feasible solution for the transformed optimization problem and, thus, the second inequality holds.

A.1.2.3 Proof: $V_t^{\alpha}(c) \geq \tilde{V}_t^{\alpha}(c)$ for $p_{t,c}^{\alpha} \in \left[0, p_t^{\alpha,\infty}\right)$.

It holds, that $\Delta_{t,c}^{\alpha} < p_{t,c}^{\alpha} < p_t^{\alpha,\infty}$, because, if $V_{t-1}^{\alpha}(c) - V_{t-1}^{\alpha}(c-1) < p_{t,c}^{\alpha}$ would not hold, the objective value would be less than $V_{t-1}^{\alpha}(c)$ and we would be better off with the null price p_t^{∞}. Consequently, case A.1.2.1 would apply and that would contradict our assumption. From the induction hypothesis, we have $\tilde{\Delta}_{t,c}^{\alpha} < p_t^{\alpha,\infty}$. Thus, there are prices between the opportunity cost and the null price $p_t^{\alpha,\infty}$ and we have $\tilde{\Delta}_{t,c}^{\alpha} = \tilde{V}_{t-1}^{\alpha}(c) - \tilde{V}_{t-1}^{\alpha}(c-1) < \tilde{p}_{t,c}^{\alpha} < p_t^{\alpha,\infty}$.

We now continue by considering the optimal price in the transformed, risk-neutral problem $\tilde{p}_{t,c}^{\alpha}$ and the corresponding optimal objective value $\tilde{V}_t^{\alpha}(c)$. Moreover,

$$
\begin{aligned}
\tilde{V}_t^{\alpha}(c) = \frac{1 - d_t\left(\tilde{p}_{t,c}^{\alpha}\right)}{\alpha} \cdot \tilde{V}_{t-1}^{\alpha}(c) + \left(1 - \frac{1 - d_t\left(\tilde{p}_{t,c}^{\alpha}\right)}{\alpha}\right) \cdot \left(\tilde{p}_{t,c}^{\alpha} + \tilde{V}_{t-1}^{\alpha}(c-1)\right) \underbrace{=}_{\mathrm{IH}} \frac{1 - d_t\left(p_{t,c}^{T,\alpha}\right)}{\alpha} \\
\cdot V_{t-1}^{\alpha}(c) + \left(1 - \frac{1 - d_t\left(\tilde{p}_{t,c}^{\alpha}\right)}{\alpha}\right) \cdot \left(\tilde{p}_{t,c}^{\alpha} + V_{t-1}^{\alpha}(c-1)\right) \\
= \min_{\substack{z_{t-1,c-1}, \\ z_{t-1,c}}} \left(1 - d_t\left(\tilde{p}_{t,c}^{\alpha}\right)\right) \cdot z_{t-1,c} \cdot V_{t-1}^{\alpha}(c) + d_t\left(\tilde{p}_{t,c}^{\alpha}\right) \cdot z_{t-1,c-1} \\
\cdot \left(\tilde{p}_{t,c}^{\alpha} + V_{t-1}^{\alpha}(c-1)\right) \leq V_t^{\alpha}(c)
\end{aligned}
$$

The third equality can be explained as follows. From $\Delta_{t,c}^{\alpha} < \tilde{p}_{t,c}^{\alpha} < p_t^{\alpha,\infty}$ (see above), it follows that $z_{t-1,c} = \frac{1}{\alpha}$ and $z_{t-1,c-1} = \left(1 - \frac{1 - d_t\left(\tilde{p}_{t,c}^{\alpha}\right)}{\alpha}\right) / d_t\left(\tilde{p}_{t,c}^{\alpha}\right)$ is the optimal solution for the minimization. The inequality follows by the fact that $\tilde{p}_{t,c}^{\alpha}$ (together with $z_{t-1,c}$ and $z_{t-1,c-1}$) is a feasible solution for the risk-averse problem.

A.1.2.4 Conclusion of A.1.2.2 and A.1.2.3

We first consider equivalence of the value functions. In the case covered in Section A.1.2.1, we directly showed the equivalence. In Section A.1.2.2, we saw that $p_{t,c}^{\alpha}$ is a feasible solution for the transformed risk-neutral optimization problem with an objective value that is greater than or equal to $V_t^{\alpha}(c)$. In Section A.1.2.3, we have shown that $\tilde{p}_{t,c}^{\alpha}$ is a feasible solution for the risk-averse optimization problem with an objective value that is greater than or equal to $\tilde{V}_t^{\alpha}(c)$. Hence, $V_t^{\alpha}(c) = \tilde{V}_t^{\alpha}(c)$. Moreover, it holds that every optimal solution $p_{t,c}^{\alpha} \in \left[0, p_t^{\alpha,\infty}\right)$ is also an optimal solution for the transformed risk-neutral optimization problem (and vice versa) as it is feasible and leads to the optimal objective value. As the transformed problem is a standard dynamic pricing problem, we know that it has a unique solution, namely $\tilde{p}^{\alpha}_{t,c}$. Consequently it holds that $p^{\alpha}_{t,c} = \tilde{p}^{\alpha}_{t,c}$ on the interval $[0, p^{\alpha,\infty}_t)$, and hence, the risk-averse optimization problem has also a unique solution on this interval.

A.2 Proof of Lemma 2

First, we consider $d_t\left(\Delta^{\alpha}_{t,c}\right) > 1 - \alpha$. As the selling probability is decreasing in the price, the non-emptiness of the interval $\left(\Delta^{\alpha}_{t,c}, p^{\alpha,\infty}_t\right)$ is a direct consequence. Together with Theorem 1, it follows that $\tilde{V}^{\alpha}_{t-1}(c) - \tilde{V}^{\alpha}_{t-1}(c-1) < p^{\alpha,\infty}_t$. Next, note that the optimal price is never below opportunity cost. If so, selling had lower total revenue than not selling. A higher price (up to the opportunity cost) would simultaneously decrease the probability of this undesired event and improve its outcome. Thus, setting a price lower than the opportunity costs cannot be optimal (see Hempenius 1970, p. 4 for a formal proof and the markup on costs). Accordingly, the optimal price of the transformed risk-neutral optimization problem $\tilde{p}^{\alpha}_{t,c}$ is an element of the non-empty interval $\left(\tilde{\Delta}^{\alpha}_{t,c}, p^{\alpha,\infty}_t\right)$, and because of $p^{\alpha}_{t,c} = \tilde{p}^{\alpha}_{t,c}$, we have $p^{\alpha}_{t,c} \in \left(\Delta^{\alpha}_{t,c}, p^{\alpha,\infty}_t\right)$.

Next, we consider $d_t\left(\Delta^{\alpha}_{t,c}\right) \leq 1 - \alpha$. As $p^{\alpha,\infty}_t$ was defined such that $d_t(p^{\alpha,\infty}_t) = 1 - \alpha$ and $d_t(p)$ is decreasing, we have $\Delta^{\alpha}_{t,c} \geq p^{\alpha,\infty}_t$. Because setting prices below $\Delta^{\alpha}_{t,c}$ is never optimal (see above), the optimal price in the transformed problem is $\tilde{p}^{\alpha}_{t,c} = p^{\alpha,\infty}_t$ with the transformed selling probability $\tilde{d}^{\alpha}_t(p^{\alpha,\infty}_t) = 0$. From Theorem 1, we have $p^{\alpha}_{t,c} = p^{\infty}_t$.

A.3 Proof of Lemma 3

To prove the statement of Lemma 3, we use Theorem 1 to transform the risk-averse dynamic pricing problem to a risk-neutral dynamic pricing problem. Showing $d_t\left(\Delta^{\alpha}_{t,c}\right) > 1 - \alpha$ is equivalent to showing $\Delta^{\alpha}_{t,c} < p^{\alpha,\infty}_t$ because $d_t(p)$ is decreasing in p and, by definition, $d_t(p^{\alpha,\infty}_t) = 1 - \alpha$. For $p^{\alpha,\infty}_t = \infty$, $\Delta^{\alpha}_{t,c} < p^{\alpha,\infty}_t$ is obvious. If $p^{\alpha,\infty}_t$ is finite, the inequality holds with:

$$\Delta^{\alpha}_{t,c} = V^{\alpha}_t(c) - V^{\alpha}_t(c-1) = \tilde{V}^{\alpha}_t(c) - \tilde{V}^{\alpha}_t(c-1) \leq \tilde{V}^{\alpha}_t(c) - \tilde{V}'^{\,\alpha}_t(c-1) < p^{\alpha,\infty}_t$$

This can be explained as follows. By definition and Theorem 1, we have $\Delta_{t,c}^{\alpha} = V_t^{\alpha}(c) - V_t^{\alpha}(c-1) = \tilde{V}_t^{\alpha}(c) - \tilde{V}_t^{\alpha}(c-1)$. Slightly abusing notation, we now denote the optimal policy from state (t,c) onwards by $\overline{\tilde{p}^{\alpha}}$. This policy is leading to the value $\tilde{V}_t^{\alpha}(c)$. There is also an optimal policy from state $(t, c-1)$ onwards that is leading to $\tilde{V}_t^{\alpha}(c-1)$. But instead of using this policy, we use a policy $\overline{\tilde{p}^{\alpha}}'$ that copies to some degree the optimal policy from state (t,c) onwards. If capacity is left, this new policy just uses the former's price. If no capacity is left (that is, if $c-1$ products are sold during the evolution of the selling process), the null price $p_t^{\alpha,\infty}$ is used to ensure that $\tilde{d}_t^{\alpha}(p_t^{\alpha,\infty}) = 0$. This new policy $\overline{\tilde{p}^{\alpha}}'$ is formally defined as follows:

$$\overline{\tilde{p}_{t,c}^{\alpha}}' = \begin{cases} \overline{\tilde{p}_{t,c}^{\alpha}}, & c > 1 \\ p_t^{\alpha,\infty}, & c = 1 \end{cases} \quad \forall\, t, c$$

From this new policy $\overline{\tilde{p}^{\alpha}}'$, we obtain an expected value of $\tilde{V}_t'^{\,\alpha}(c-1)$. As this policy is not necessarily optimal, we have $\tilde{V}_t'^{\,\alpha}(c-1) \leq \tilde{V}_t^{\alpha}(c-1)$.

We now compare the evolution of the selling process starting in (t,c) following policy $\overline{\tilde{p}^{\alpha}}'$ and starting in $(t, c-1)$ following policy $\overline{\tilde{p}^{\alpha}}'$. As the latter copies the first policy (to some degree), the evolution of the selling process is the same at the beginning. It holds that for every realization of demand throughout the selling horizon without sellout of $c-1$ products (i.e. if strictly less than $c-1$ customers buy), the earned revenue and the corresponding probability of occurrence are the same for the processes starting in (t,c) and $(t,c-1)$. The only difference is if $c-1$ products are sold and there is still time left to sell the last product, which is only available in the process that started at (t,c). The probability of occurrence of this situation is below 1 and the highest revenue that can be earned by selling the last product is strictly less than $p_t^{\alpha,\infty}$. Thus, we have $\tilde{V}_t^{\alpha}(c) - \tilde{V}_t'^{\,\alpha}(c-1) < p_t^{\alpha,\infty}$.

A.4 Proof of Theorem 2

In this proof, we show by induction that the optimal objective values of both optimization problems are equal and, simultaneously, that the optimal risk-averse price $p_{t,c}^{\alpha,\lambda}$ equals the optimal price $\tilde{p}_{t,c}^{\alpha,\lambda}$ in the transformed risk-neutral problem (and vice versa). For $t = 0$, the boundary conditions of both optimization problems are the same and, thus, the objective values are the same (induction basis). In the induction step, we assume that $V_{t-1}^{\alpha,\lambda}(c) = \tilde{V}_{t-1}^{\alpha,\lambda}(c)\ \forall\ c$ (induction hypothesis, IH) and show that this equality also holds for t. Simultaneously, we confirm that $p_{t,c}^{\alpha,\lambda} = \tilde{p}_{t,c}^{\alpha,\lambda}\quad \forall\ t, c$.

For the induction step, we rewrite both value functions as follows: $V_t^{\alpha,\lambda}(c) = \max_p h_{t,c}^{\alpha,\lambda}(p)$ with

$$h_{t,c}^{\alpha,\lambda}(p) = (1-\lambda)\mathbb{E}\left[1_{\{X_t \geq p\}} \cdot p + V_{t-1}^{\alpha,\lambda}\left(c - 1_{\{X_t \geq p\}}\right)\right] + \lambda\,\mathrm{CVaR}_\alpha\left(1_{\{X_t \geq p\}} \cdot p + V_{t-1}^{\alpha,\lambda}\left(c - 1_{\{X_t \geq p\}}\right)\right)$$

and $\tilde{V}_t^{\alpha,\lambda}(c) = \max_p \tilde{h}_{t,c}^{\alpha,\lambda}(p)$ with $\tilde{h}_{t,c}^{\alpha,\lambda}(p) = \tilde{\mathbb{E}}^{\alpha,\lambda}\left[1_{\{X_t \geq p\}} \cdot p + \tilde{V}_{t-1}^{\alpha,\lambda}\left(c - 1_{\{X_t \geq p\}}\right)\right]$ where $\tilde{\mathbb{E}}^{\alpha,\lambda}$ denotes the expectation with regard to the modified selling probability $\tilde{d}_t^{\alpha,\lambda}(p)$.

In both value functions, the minimization is over all $p \in \mathcal{P}_t = [0, p_t^\infty]$. However, it suffices to show that $h_{t,c}^{\alpha,\lambda}(p) = \tilde{h}_{t,c}^{\alpha,\lambda}(p)\ \forall t, c, \Delta_t^{\alpha,\lambda}(c) \underbrace{=}_{\mathrm{IH}} \tilde{\Delta}_{t,c}^{\alpha,\lambda} \leq p \leq p_t^\infty$, given the induction hypothesis holds for $t-1$. For prices below opportunity cost ($p < \Delta_t^{\alpha,\lambda}(c) \underbrace{=}_{\mathrm{IH}} \tilde{\Delta}_{t,c}^{\alpha,\lambda}$), remember that on the one hand, these prices are never optimal. Now, selling is the undesired event and CVaR as well as expected value improve by increasing the price, which improves the undesired event and decreases its probability. Thus, we have $h_{t,c}^{\alpha,\lambda}(p) < h_{t,c}^{\alpha,\lambda}\left(\Delta_{t,c}^{\alpha,\lambda}\right)$ for $\forall\ p < \Delta_{t,c}^{\alpha,\lambda}$. On the other hand, we have $\tilde{h}_{t,c}^{\alpha,\lambda}(p) < \tilde{h}_{t,c}^{\alpha,\lambda}\left(\tilde{\Delta}_{t,c}^{\alpha,\lambda}\right)\ \forall\ p < \tilde{\Delta}_{t,c}^{\alpha,\lambda}$ because this is a standard dynamic pricing problem where, again, selling below opportunity cost makes no sense. Accordingly, the maximizing p will not be below opportunity cost.

Thus, in the following, we only consider $\Delta_t^{\alpha,\lambda}(c) \underbrace{=}_{\mathrm{IH}} \tilde{\Delta}_{t,c}^{\alpha,\lambda} \leq p \leq p_t^\infty$. Accordingly, we explicitly solve CVaR's minimization using the fact that selling (as opposed to not selling) is never the worse event. We distinguish two cases for $p \in \left[\Delta_t^{\alpha,\lambda}(c) = \tilde{\Delta}_{t,c}^{\alpha,\lambda},\ p_t^\infty\right]$. Moreover, note that if opportunity costs exceed the null price p_t^∞, we select it and $h_{t,c}^{\alpha,\lambda}(p_t^\infty) = \tilde{h}_{t,c}^{\alpha,\lambda}(p_t^\infty)$ holds analogous to case A.4.1.

A.4.1 $h_{t,c}^{\alpha,\lambda}(p) = \tilde{h}_{t,c}^{\alpha,\lambda}(p)$, $0 \leq d_t(p) < 1-\alpha$

In this case, the selling probability is so low that CVaR includes only the event of not selling and we have $\mathrm{CVaR}_\alpha\left(1_{\{X_t \geq p\}} \cdot p + V_{t-1}^{\alpha,\lambda}\left(c - 1_{\{X_t \geq p\}}\right)\right) = V_{t-1}^{\alpha,\lambda}(c)$. We have

$$h_{t,c}^{\alpha,\lambda}(p) = (1-\lambda)\mathbb{E}\left[1_{\{X_t \geq p\}} \cdot p + V_{t-1}^{\alpha,\lambda}\left(c - 1_{\{X_t \geq p\}}\right)\right] + \lambda V_{t-1}^{\alpha,\lambda}(c)$$

$$= (1-\lambda)\left[\left(1 - d_t(p)\right) \cdot V_{t-1}^{\alpha,\lambda}(c) + d_t(p) \cdot \left(p + V_{t-1}^{\alpha,\lambda}(c-1)\right)\right] + \lambda V_{t-1}^{\alpha,\lambda}(c)$$

$$\underbrace{=}_{\text{IH}} \left[(1-\lambda)\left(1 - d_t(p)\right) + \lambda\right] \cdot \tilde{V}_{t-1}^{\alpha,\lambda}(c) + (1-\lambda) d_t(p) \cdot \left(p + \tilde{V}_{t-1}^{\alpha,\lambda}(c-1)\right)$$

$$= \left[1 - (1-\lambda) d_t(p)\right] \cdot \tilde{V}_{t-1}^{\alpha,\lambda}(c) + (1-\lambda) d_t(p) \cdot \left(p + \tilde{V}_{t-1}^{\alpha,\lambda}(c-1)\right)$$

$$= \left[1 - d_t^{T,\alpha,\lambda}(p)\right]\left(\tilde{V}_{t-1}^{\alpha,\lambda}(c)\right) + d_t^{T,\alpha,\lambda}(p) \cdot \left(p + \tilde{V}_{t-1}^{\alpha,\lambda}(c-1)\right)$$

$$= \tilde{\mathbb{E}}^{\alpha,\lambda}\left[1_{\{X_t \geq p\}} \cdot p + \tilde{V}_{t-1}^{\alpha,\lambda}\left(c - 1_{\{X_t \geq p\}}\right)\right] = \tilde{h}_{t,c}^{\alpha,\lambda}(p).$$

A.4.2 $h_{t,c}^{\alpha,\lambda}(p) = \tilde{h}_{t,c}^{\alpha,\lambda}(p)$, $d_t(p) \geq 1 - \alpha$

In this case, the selling probability is so high that CVaR includes both the events of selling and not selling and we have

$$\mathrm{CVaR}_\alpha\left(1_{\{X_t \geq p\}} \cdot p + V_{t-1}^{\alpha,\lambda}\left(c - 1_{\{X_t \geq p\}}\right)\right)$$

$$= \min_{\substack{z_{t-1,c-1},\\ z_{t-1,c}}} \left(1 - d_t(p)\right) \cdot z_{t-1,c} \cdot V_{t-1}^{\alpha,\lambda}(c) + d_t(p) \cdot z_{t-1,c-1} \cdot \left(p + V_{t-1}^{\alpha,\lambda}(c-1)\right)$$

$$= \left(1 - d_t(p)\right) \cdot \frac{1}{\alpha} \cdot V_{t-1}^{\alpha,\lambda}(c) + d_t(p) \cdot \left(1 - \frac{1 - d_t(p)}{\alpha}\right) / d_t\left(p_{t,c}^{\alpha,\lambda}\right) \cdot \left(p + V_{t-1}^{\alpha,\lambda}(c-1)\right)$$

$$= \left(1 - d_t(p)\right) \cdot \frac{1}{\alpha} \cdot V_{t-1}^{\alpha,\lambda}(c) + \left(1 - \frac{1 - d_t(p)}{\alpha}\right) \cdot \left(p + V_{t-1}^{\alpha,\lambda}(c-1)\right)$$

where the second equality again uses the fact that selling is never the worse event and, thus, $z_{t-1,c-1} = \left(1 - \frac{1-d_t(p)}{\alpha}\right) / d_t(p)$ and $z_{t-1,c} = 1/\alpha$ are optimal.

Thus, we have

$$h_{t,c}^{\alpha,\lambda}(p) = (1-\lambda)\mathbb{E}\left[1_{\{X_t \geq p\}} \cdot p_{t,c}^{\alpha,\lambda} + V_{t-1}^{\alpha,\lambda}\left(c - 1_{\{X_t \geq p\}}\right)\right]$$

$$+\lambda\left[\left(1 - d_t(p)\right) \cdot \frac{1}{\alpha} \cdot V_{t-1}^{\alpha,\lambda}(c) + \left(1 - \frac{1 - d_t(p)}{\alpha}\right) \cdot \left(p + V_{t-1}^{\alpha,\lambda}(c-1)\right)\right]$$

$$= (1-\lambda)\left[\left(1 - d_t(p)\right) \cdot V_{t-1}^{\alpha,\lambda}(c) + d_t(p) \cdot \left(p + V_{t-1}^{\alpha,\lambda}(c-1)\right)\right]$$

$$+\lambda\left[\left(1-d_t(p)\right)\cdot\tfrac{1}{\alpha}\cdot V_{t-1}^{\alpha,\lambda}(c)+\left(1-\tfrac{1-d_t(p)}{\alpha}\right)\cdot\left(p+V_{t-1}^{\alpha,\lambda}(c-1)\right)\right]$$

$$\underbrace{=}_{\text{IH}}\left[(1-\lambda)\left(1-d_t(p)\right)+\lambda\left(1-d_t(p)\right)\cdot\tfrac{1}{\alpha}\right]\cdot\tilde{V}_{t-1}^{\alpha,\lambda}(c)$$

$$+\left[(1-\lambda)d_t(p)+\lambda\left(1-\tfrac{1-d_t(p)}{\alpha}\right)\right]\cdot\left(p+\tilde{V}_{t-1}^{\alpha,\lambda}(c-1)\right)$$

$$=\left[1-(1-\lambda)d_t(p)-\lambda\left(1-\tfrac{1-d_t(p)}{\alpha}\right)\right]\cdot\tilde{V}_{t-1}^{\alpha,\lambda}(c)$$

$$+\left[(1-\lambda)d_t(p)+\lambda\left(1-\tfrac{1-d_t(p)}{\alpha}\right)\right]\cdot\left(p+\tilde{V}_{t-1}^{\alpha,\lambda}(c-1)\right)$$

$$=\left(1-\tilde{d}_t^{\alpha,\lambda}(p)\right)\left(\tilde{V}_{t-1}^{\alpha,\lambda}(c)\right)+\tilde{d}_t^{\alpha,\lambda}(p)\cdot\left(p+\tilde{V}_{t-1}^{\alpha,\lambda}(c-1)\right)$$

$$=\tilde{\mathbb{E}}^{\alpha,\lambda}\left[1_{\{X_t\geq p\}}\cdot p+\tilde{V}_{t-1}^{\alpha,\lambda}\left(c-1_{\{X_t\geq p\}}\right)\right]=\tilde{h}_{t,c}^{\alpha,\lambda}(p)$$

A.5 Proof of Proposition 3

The conditions of Proposition 3 are met, i.e. there is a $y\in\mathbb{R}$ such that $F_t(\alpha^y p)=\alpha\cdot F_t(p)$ and $f_t(\alpha^y p)=\alpha^{1-y}\cdot f_t(p)\ \forall\, t$. Consequently, it holds that $d_t(\alpha^y p)=1-F_t(\alpha^y p)=1-\alpha\cdot F_t(p)=1-\alpha\cdot\left(1-d_t(p)\right)=1-\alpha+\alpha\cdot d_t(p)$ and, for the derivative, $d_t'(\alpha^y p)=-f_t(\alpha^y p)=-\alpha^{1-y}\cdot f_t(p)=\alpha^{1-y}\cdot d_t'(p)$. Let $p_{t,c}^1$ denote the optimal price that a risk-neutral firm sets and $V_t^1(c)$ the corresponding value function. In the following, we will repeatedly use the f.o.c. of the problems. As they are similar, we only restate the risk-neutral f.o.c. to improve readability: $d_t'\left(p_{t,c}^1\right)\cdot\left(p_{t,c}^1-\Delta_{t,c}^1\right)+d_t\left(p_{t,c}^1\right)=0$.

We now show the proposition by induction over t, i.e. that $p_{t,c}=\alpha^y\cdot p_{t,c}^1$ and $V_t^\alpha(c)=\alpha^y\cdot V_t^1(c)$.

For $t=1$, we plug $\alpha^y\cdot p_{1,c}^1$ in the sufficient condition (the f.o.c. of the transformed problem) and derive

$$0=\left(\tilde{d}_1^\alpha\right)'\left(\alpha^y\cdot p_{1,c}^1\right)\cdot\alpha^y\cdot p_{1,c}^1+\tilde{d}_1^\alpha\left(\alpha^y\cdot p_{1,c}^1\right)=\frac{1}{\alpha}d_1'\left(\alpha^y\cdot p_{1,c}^1\right)\cdot\alpha^y\cdot p_{1,c}^1+1-\frac{1-d_1\left(\alpha^y\cdot p_{1,c}^1\right)}{\alpha}$$

$$=\alpha^{-y}\cdot d_1'\left(p_{1,c}^1\right)\cdot\alpha^y\cdot p_{1,c}^1+1-\frac{\alpha-\alpha\cdot d_1\left(p_{1,c}^1\right)}{\alpha}=d_1'\left(p_{1,c}^1\right)\cdot p_{1,c}^1+d_1\left(p_{1,c}^1\right)$$

where the first equality follows by definition of $\tilde{d}_1^{\alpha}(p)$ (Theorem 1). The second follows from $d_t(\alpha^y p)$ and $d_t{}'(\alpha^y p)$ discussed above and the third is a rearrangement. The fourth shows that the f.o.c. of the transformed problem is equivalent to the f.o.c. of the risk-neutral problem. Obviously, these equations imply that $\alpha^y \cdot p_{1,c}^1$ is an optimal solution to the risk-averse problem if and only if $p_{1,c}^1$ is an optimal solution to the risk-neutral problem.

Next, by using the same considerations as above, we calculate $V_1^{\alpha}(c) = \tilde{d}_1^{\alpha}\left(p_{1,c}^{\alpha}\right) \cdot p_{1,c}^{\alpha} = \tilde{d}_1^{\alpha}\left(\alpha^y \cdot p_{1,c}^1\right) \cdot \alpha^y \cdot p_{1,c}^1 = \left(1 - \frac{1-d_1\left(\alpha^y \cdot p_{1,c}^1\right)}{\alpha}\right) \cdot \alpha^y \cdot p_{1,c}^1 = d_1\left(p_{1,c}^1\right) \cdot \alpha^y \cdot p_{1,c}^1 = \alpha^y \cdot V_1^1(c)$. Thus, the proposition holds for $t = 1$.

In the induction step, we assume that $V_{t-1}^{\alpha}(c) = \alpha^y \cdot V_{t-1}^1(c)$ (induction hypothesis, IH) and show that this equality also holds for t. By doing so, we also show that $p_{t,c}^{\alpha} = \alpha^y \cdot p_{t,c}^1$. As $V_{t-1}^{\alpha}(c) = \alpha^y \cdot V_{t-1}^1(c)$ (IH), it holds that $\Delta_{t,c}^{\alpha} = \alpha^y \cdot \Delta_{t,c}^1$ (with $\Delta_{t,c}^1 = V_{t-1}^1(c) - V_{t-1}^1(c-1)$). To do so, we now distinguish two cases:

A.5.1 Case 1: $\Delta_{t,c}^{1} < p_t^{\infty} \Leftrightarrow \tilde{\Delta}_{t,c}^{\alpha} < p_t^{\alpha,\infty}$

In this case, opportunity costs are below the null price and it makes sense to try to sell a unit. It is important so see that this happens simultaneously in both problems – the risk-neutral one and the transformed, risk-averse problem. More technically, to show this equivalence of the two conditions, note that $\tilde{d}_1^{\alpha}\left(\tilde{\Delta}_{t,c}^{\alpha}\right) = \tilde{d}_1^{\alpha}\left(\alpha^y \cdot \Delta_{t,c}^1\right) = 1 - \left(1 - d_t\left(\alpha^y \cdot \Delta_{t,c}^1\right)\right) / \alpha = 1 - \left(1 - \left(1 - \alpha + \alpha \cdot d_t\left(\Delta_{t,c}^1\right)\right)\right) / \alpha = d_t\left(\Delta_{t,c}^1\right)$. Using this, the second condition can be reformulated as $\tilde{\Delta}_{t,c}^{\alpha} < p_t^{\alpha,\infty} \Leftrightarrow d_t\left(\Delta_{t,c}^1\right) = \tilde{d}_1^{\alpha}\left(\tilde{\Delta}_{t,c}^{\alpha}\right) > \tilde{d}_1^{\alpha}\left(p_t^{\alpha,\infty}\right) = 1 - \left(1 - \tilde{d}_1^{\alpha}\left(p_t^{\alpha,\infty}\right)\right) / \alpha = 1 - \left(1 - (1-\alpha)\right) / \alpha = 0$, where the second inequality holds because $\tilde{d}_1^{\alpha}(\cdot)$ is strictly decreasing. From the first condition, we have $\Delta_{t,c}^1 < p_t^{\infty} \Leftrightarrow d_t\left(\Delta_{t,c}^1\right) > d_t\left(p_t^{\alpha,\infty}\right)$, as $d_t'(\cdot) < 0$.

Now, if $p_t^{\infty} = +\infty$, then the first condition is always satisfied and the second also because no matter how high $\Delta_{t,c}^1$, $d_t\left(\Delta_{t,c}^1\right)$ is always positive. Thus, the conditions are equivalent.

If p_t^{∞} is finite, we have $d_t(p_t^{\infty}) = 0$. Again, both conditions are equivalent.

Similar to $t = 1$, we again show that $\alpha^y \cdot p_{t,c}^1$ is an optimal solution to the risk-averse problem if and only if $p_{t,c}^1$ is an optimal solution to the risk-neutral problem.

From f.o.c., we have $d_t'(p_{t,c}^1) \cdot (p_{t,c}^1 - \Delta_{t,c}^1) + d_t(p_{t,c}^1) = 0$. Using Theorem 1, we know that $p_{t,c}^\alpha$ (the optimal price a risk-averse firm sets) fulfills the f.o.c. $(\tilde{d}_1^\alpha)'(p) \cdot (p - \Delta_{t,c}^\alpha) + \tilde{d}_1^\alpha(p) = 0$ (as the transformed optimization problem described in Theorem 1 with $\tilde{d}_1^\alpha(p) = 1 - (1 - d_t(p))/\alpha$ is equivalent to the risk-averse optimization problem). We plug $\alpha^y \cdot p_{t,c}^1$ in the risk-averse f.o.c. and derive

$$\begin{aligned}
&(\tilde{d}_1^\alpha)'(\alpha^y \cdot p_{t,c}^1) \cdot (\alpha^y \cdot p_{t,c}^1 - \Delta_{t,c}^\alpha) + \tilde{d}_1^\alpha(\alpha^y \cdot p_{t,c}^1) \\
&\qquad = \frac{1}{\alpha} d_t'(\alpha^y \cdot p_{t,c}^1) \cdot (\alpha^y \cdot p_{t,c}^1 - \alpha^y \cdot \Delta_{t,c}^1) + 1 - \frac{1 - d_t(\alpha^y \cdot p_{t,c}^1)}{\alpha} \\
&\qquad = \alpha^{-y} \cdot d_t'(p_{t,c}^1) \cdot \alpha^y \cdot (p_{t,c}^1 - \Delta_{t,c}^1) + 1 - \frac{\alpha - \alpha \cdot d_t(p_{t,c}^1)}{\alpha} \\
&\qquad = d_t'(p_{t,c}^1) \cdot (p_{t,c}^1 - \Delta_{t,c}^1) + d_t(p_{t,c}^1) = 0.
\end{aligned}$$

This shows that $\alpha^y \cdot p_{t,c}^1$ is a solution to the risk-averse f.o.c. if and only if $p_{t,c}^1$ solves the risk-neutral f.o.c. Again, we calculate $V_t^\alpha(c) = \tilde{d}_1^\alpha(p_{t,c}^\alpha) \cdot (p_{t,c}^\alpha - \Delta_{t,c}^\alpha) + V_{t-1}^\alpha(c) = \tilde{d}_1^\alpha(\alpha^y \cdot p_{t,c}^1) \cdot (\alpha^y \cdot p_{t,c}^1 - \alpha^y \cdot \Delta_{t,c}^1) + \alpha^y \cdot V_{t-1}^1(c) = \left(1 - \frac{1 - d_t(\alpha^y \cdot p_{t,c}^1)}{\alpha}\right) \cdot (\alpha^y \cdot p_{t,c}^1 - \alpha^y \cdot \Delta_{t,c}^1) + \alpha^y \cdot V_{t-1}^1(c) = d_t(p_{t,c}^1) \cdot \alpha^y \cdot (p_{t,c}^1 - \Delta_{t,c}^1) + \alpha^y \cdot V_{t-1}^1(c) = \alpha^y \cdot V_t^1(c)$.

A.5.2 Case 2: $\Delta_{t,c}^{1} \geq p_t^\infty \Leftrightarrow \tilde{\Delta}_{t,c}^{\alpha} \geq p_t^{\alpha,\infty}$

In this case, opportunity costs exceed the null price and it does not make sense to try to sell a unit. From the proof of Lemma 2 (A.2), we know that selling at a price below opportunity cost is never optimal. Thus, selling in period t is avoided by setting the risk-neutral (p_t^∞) or risk-averse ($p_t^{\alpha,\infty}$) null price.