

Dynamic Pricing through Data Sampling

Maxime C. Cohen

NYU Stern School of Business, New York City, New York 10012, USA, maxcohen@nyu.edu

Ruben Lobel

Airbnb, San Francisco, California 94103, USA, ruben.lobel@gmail.com

Georgia Perakis*

Sloan School of Management, MIT, Cambridge, Massachusetts 02139, USA, georgiap@mit.edu

We study a dynamic pricing problem, where a firm offers a product to be sold over a fixed time horizon. The firm has a given initial inventory level, but there is uncertainty about the demand for the product in each time period. The objective of the firm is to determine a dynamic pricing strategy that maximizes revenue throughout the entire selling season. We develop a tractable optimization model that directly uses demand data, therefore creating a practical decision tool. We show computationally that regret-based objectives can perform well when compared to average revenue maximization and to a Bayesian approach. The modeling approach proposed in this study could be particularly useful for risk-averse managers with limited access to historical data or information about the true demand distribution. Finally, we provide theoretical performance guarantees for this sampling-based solution.

Key words: dynamic pricing; data-driven; sampling-based optimization

History: Received: April 2015; Accepted: January 2018 by William L. Cooper, after 2 revisions.

1. Introduction

In many industries, managers are faced with the challenge of selling a fixed amount of inventory within a specific time horizon. In the classical dynamic pricing problem, the firm's goal is to set prices at each stage of the selling season so as to maximize the total revenue while facing an uncertain demand. As in most real-life applications, the intrinsic randomness of the firm's demand is an important factor that should be taken into account. In some applications, such as apparel retail stores, selling products for the current season typically involves making decisions with very limited demand information. Decisions based on deterministic demand forecasts can expose the firm to severe revenue losses. The stochastic optimization framework uses knowledge of the demand distribution in order to improve the future pricing decisions. In practice, this approach can sometimes lead to very large data requirements. It also usually assumes that the demand uncertainty is independent across time periods, otherwise raising tractability issues. In many cases, however, firms have limited information about the demand distribution, and therefore cannot apply stochastic optimization. With this in mind, we propose an optimization model that directly use the available demand data, creating a practical analytic tool for pricing problems.

This distribution-free modeling approach is the motivation behind the robust pricing literature, where the demand shocks are assumed to lie in a (bounded) uncertainty set. The goal in this case is to find a pricing policy that robustly maximizes the revenue within this uncertainty set, without further assumptions on the distribution or on the demand correlation across time periods. Nevertheless, as a drawback, the robust solution is often regarded as too conservative and the robust optimization literature has mainly focused on static problems (open-loop policies). The dynamic models that search for closed-loop policies, that is, policies that adjust the pricing decisions depending on previously realized uncertainties, can often become intractable. Our goal in this study is to propose an approach that can tackle these two issues, by developing a framework for finding non-conservative and adjustable robust pricing policies.

1.1. Contributions

In this study, we propose a data-driven decision tool for solving the robust dynamic pricing problem. This tool is both flexible and easy to implement. It is particularly useful in retail settings where managers are risk-averse and each selling season is different, limiting the available demand information. By directly using data samples in the optimization formulation, we convexify the problem while allowing for a wide

range of modeling choices. More specifically, we use regret-based objectives and linear control policies that address two problems commonly found in the robust optimization literature: conservative objective and open-loop policies.

While the main contribution of this study lies in the applied domain, we also provide two theoretical contributions to the existing literature: (i) We prove that the sampled problem optimal solution converges (almost surely) to the original robust problem, and (ii) We extend the sampling size bound of Calafiore and Campi (2006) to admit sampling from an artificial distribution. In addition, we study and interpret some interesting properties of the sampled dual problem. Finally, we computationally demonstrate that the regret-based robust models perform well in practice, when compared to a sample-average benchmark and to a Bayesian approach. In particular, our models can be very useful when dealing with small sample sizes or when the sampling distribution is different from the true distribution.

1.2. Literature Review

A good introduction to the field of revenue management and dynamic pricing can be found in the overview paper of Elmaghraby and Keskinocak (2003) and in the book by Talluri and Van Ryzin (2006). Most of the pricing literature uses the stochastic optimization framework, which relies on distributional assumptions on the demand model and often does not capture the correlation of the demand uncertainty across time periods. To avoid such problems, there are two different prominent approaches in the literature that will be most relevant to our research: data-driven and robust optimization. For the most part, these approaches have been studied separately and the modeling choice usually depends on the type of information provided to the firm about demand uncertainty.

The operations management literature has explored sampling-based optimization as a form of data-driven non-parametric approach for solving stochastic optimization problems with unknown distributions. In this case, it is common to use past historical data, which are sample evaluations coming from the true distribution of the uncertain parameter. A typical form of data-driven approach is known as the sample average approximation (SAA), in which the scenario evaluations are averaged to approximate the expectation of the objective function. Kleywegt et al. (2001) deal with a discrete stochastic optimization model and show that the SAA solution converges almost surely to the optimal solution of the original problem when the number of samples goes to infinity. The authors also derive a bound on the number of samples required to obtain at most a certain difference between the SAA solution and the optimal value,

under some confidence level. In the dynamic pricing literature, the data-driven approach has been used by Rusmevichientong et al. (2006) to develop a non-parametric data-driven approach to pricing, and also more recently by Eren and Maglaras (2010). Bertsimas and Vayanos (2017) propose a pricing model with learning that captures the exploration vs. exploitation trade-off without distributional assumptions which are common in the learning literature.

The field of robust optimization studies distribution-free modeling ideas for making decisions under uncertainty (see, e.g., Ben-Tal and Nemirovski 2000, Bertsimas and Sim 2004). A robust policy can be defined in several ways. In this study, we will explore three different types of robust models: the MaxMin, the MinMax Regret (or alternatively MinMax Absolute Regret) and the MaxMin Ratio (or alternatively MaxMin Relative Regret or MaxMin Competitive Ratio). In inventory management, the MaxMin robust approach can be seen for example in Gallego and Moon (1993) and Bertsimas and Thiele (2006). The paper by Adida and Perakis (2006) is an example of the MaxMin robust approach applied to the dynamic pricing problem. This approach is usually appropriate for risk-averse managers, but can often yield conservative solutions. For this reason, we will also explore the regret-based models. Lim et al. (2009) approach this problem from a different angle, where the pricing policies are hedged against a family of distributions bounded by a relative entropy measure. A comparison of the MaxMin and the MinMax Regret objectives in revenue management can be found in Perakis and Roels (2010). An alternative approach is the relative regret measure, also known as the competitive ratio. In revenue management and pricing, Ball and Queyranne (2009) use such a MaxMin Ratio approach.

An additional common methodology used in dynamic pricing takes a Bayesian approach. More precisely, a parametric demand model is postulated jointly with a prior distribution that reflects the seller's initial knowledge of the parameters. Demand observations are then collected and used to update the prior into a posterior distribution via Bayes' rule. Examples of such works include Aviv and Pazgal (2002), Lin (2006), Araman and Caldentey (2009) and Farias and Van Roy (2010).

Irrespective of the optimization objectives described above, multi-period decision models can be categorized as follows: (i) closed-loop policies, which use feedback from the actual state of the system at each stage, and (ii) open-loop policies, which are defined statically at the beginning of the horizon. The robust framework used in the papers above does not typically allow for adaptability in the optimal policy. This open-loop robust framework has been applied to

the dynamic pricing problem in e.g., Adida and Perakis (2006). Moving toward closed-loop solutions, Ben-Tal et al. (2004) first introduced adaptability to robust optimization problems. Ben-Tal et al. (2005) propose an application of adjustable robust optimization in a supply chain problem. More specifically, the authors advocate for the use of affinely adjustable policies. Bertsimas et al. (2010) was able to show that optimality can actually be achieved by affine policies for a particular class of one-dimensional multi-stage robust problems. Bertsimas et al. (2011) further extends this class of problems by using polynomial policies that can be solved via semi-definite programming.

As mentioned in Caramanis (2006), the sampling approach to the adaptable robust problem puts aside the non-convexities created by the influence of the realized uncertainties in the policy decisions. The natural question which arises is to identify the number of scenarios needed to reach confidence guarantees of the model's solution. To answer this question, Calafiore and Campi (2005, 2006) provided a theoretical bound on the sample size required to obtain an ϵ -level robust solution. The bound was later improved by Campi and Garatti (2008) and by Calafiore (2009), with a tightness guarantee for the class of "fully-supported" problems. Pagnoncelli et al. (2009) further suggested how to use this framework to solve chance constrained problems. Vayanos et al. (2012) also used this framework in an algorithm for constraint sampling. Contributing to this literature, our study develops a sample size bound that can be used even when sampling from an artificial distribution which is different from the true distribution.

1.2.1. Structure of the Paper. The remainder of the study is structured as follows. In section 2, we introduce our modeling approach. In section 3, we show the simulated performance of the proposed models. In section 4, we develop the theoretical performance guarantees related to convergence and sampling size. In addition, we study some properties of the sampled dual problem. Finally, in section 5, we conclude with a summary of the theoretical and numerical results. A summary of the notation used, the proofs of our theoretical results, and some additional computational tests can be found in the Appendix.

2. Model

Before introducing our general model, we motivate the problem with the following example. Suppose a firm sells a single product over a two period horizon, with a limited inventory C . Moreover, suppose the

firm has a set of N historical data samples of demand and prices for each period of the horizon. We assume for this example that the demand is a linear function of the price plus δ , which is a random noise component, that is, $\text{Demand}_t = a_t - b_t \text{Price}_t + \delta_t$. After estimating the demand function parameters (a_t, b_t) using the N data points, we are left with a set of estimation errors (i.e., the difference between the realized and estimated demand): $\delta_t^{(1)}, \dots, \delta_t^{(N)}$ for each time period t . A typical robust pricing approach would define an uncertainty set from which these errors are coming from, and choose prices that maximize the worst-case revenue scenarios within that set. First, it is not clear how one should define this uncertainty set given a pool of uncertainty samples. Second, the resulting problem can become too hard to solve, as we will show later in this section. The direct use of the uncertainty samples $\delta_t^{(i)}$ in the price optimization is what characterizes a Sampling-based optimization model, which we advocate for in this study. Our goal is to find a pricing strategy that robustly maximizes the firm's revenue with respect to the N given demand uncertainty observations:

$$\begin{aligned} \max_{p_1, p_2 \geq 0} \min_{i=1, \dots, N} & p_1(a_1 - b_1 p_1 + \delta_1^{(i)}) + p_2(a_2 - b_2 p_2 + \delta_2^{(i)}) \\ \text{s.t.} & (a_1 - b_1 p_1 + \delta_1^{(i)}) + (a_2 - b_2 p_2 + \delta_2^{(i)}) \leq C, \\ & \forall i = 1, \dots, N \end{aligned}$$

Note that the robust objective will be generalized later in this section to allow for alternative robustness measures other than the above MaxMin objective. As we mentioned, one of the main problems with the MaxMin objective is that it may be too conservative. For this reason, we will propose other types of robust objectives. Another problem in this example is that the second period price p_2 does not depend on the uncertainty realized in the first period δ_1 (open-loop policy). Ideally, the second period pricing policy should be a function of the information obtained in the first period, that is, $p_2(\delta_1)$ (a closed-loop policy).

The two main theoretical challenges with this sampling-based optimization approach are as follows: (i) Demonstrating that the solution of the sampled problem converges to the original robust problem; (ii) Finding how many samples are needed for the sampling-based model to yield a good approximation of the original problem.

We next extend the above motivating example by generalizing the problem to include components such as salvage/overbooking inventory, nonlinear demand models, adjustable pricing policies, and regret-based objectives. Throughout section 2.1, we develop a solution approach that mitigates the modeling issues

described above. We refer the reader to Appendix A for a summary of the notation used throughout the study.

2.1. Model Definition

Let T be the length of the time horizon, p_t the price at time t and d_t the nominal demand function. The nominal demand d_t captures the deterministic part of the demand at a given price level. It is worth noting that the modeling techniques and results presented in this study can be easily implemented with price effects across time periods (which also allow us to use demand models with reference prices). We are only restricting ourselves to effects of the current price to avoid complicating the notation. In addition, the model presented in this study can be extended to a multi-product pricing problem, with multiple resources. Given a multi-product demand function and a product-resource allocation matrix, both theoretical and numerical results still hold (this extension is briefly explored in Appendix C.2). Nevertheless, it should be noted that estimating a demand function with cross-product elasticities is not necessarily easy in practice. Finally, the adaptable pricing policy needs to be carefully specified in order to avoid dimensionality issues. For technical reasons, we require that the demand function satisfies the following assumption.

ASSUMPTION 1. *Let the nominal demand, $d_t(p_t)$, be a non-increasing function of the price p_t for a given set of demand parameters. We also assume that $d_t(p_t)$ is convex in p_t , and that $\sum_{t=1}^T p_t d_t(p_t)$ is strictly concave in p .*

Assumption 1 allows us to ensure that the space of pricing strategies is a closed convex set, and that the objective function is strictly concave. Consequently, this will give rise to a unique optimal solution. Several commonly used demand functions in revenue management satisfy Assumption 1, such as the linear, iso-elastic, and logarithmic functions.

In order to capture the demand variability, given the nominal demand function, $d_t(p_t)$, we define $\tilde{d}_t(p_t, \delta_t)$ as the actual demand at time t , which is realized for some uncertain parameter δ_t . For example, when using additive uncertainty, $\tilde{d}_t(p_t, \delta_t) = d_t(p_t) + \delta_t$. In the literature, it is common to use additive or multiplicative uncertainty. In our framework, we admit any sort of dependence of the demand function on the uncertain parameters.

To allow for adjustability in the pricing function, the price p_t is defined as a function of the control variables s and the uncertainty δ , thus denoted by $p_t(s, \delta)$. We also refer to s as the pricing policy. We assume that there is a finite set of decision variables s , which lie within the strategy space S . We assume that S is a finite dimensional compact set. In the case of static

pricing (open-loop policy), s is a vector of fixed prices decided upfront, independent of the demand realizations. When using adjustable policies (closed-loop), the actual price at time t must naturally be a function only of the uncertainty up to time $t - 1$. For conciseness, we express the actual realized prices as $p_t(s, \delta)$ in both cases. For a concrete example of such a pricing policy, see section 2.2, Equation (6). Also, to avoid complicating the notation, we define $\tilde{d}_t(s, \delta) = \tilde{d}_t(p_t(s, \delta), \delta_t)$. In other words, the policy s and the uncertainty δ fully determine the price at time t , and hence also the realized demand $\tilde{d}_t(s, \delta)$. In this study, we restrict ourselves to the family of pricing functions $p_t(s, \delta)$ that are affine in s . This restriction will be formalized later in Assumption 2.

In general, $\delta = (\delta_1, \dots, \delta_T)$ is a random vector with a component for each period t . We assume that δ is drawn from an unknown probability distribution \mathcal{Q} , with support on the set U , which we call the uncertainty set. We do not make any assumptions on the independence of δ across time, as opposed to most stochastic optimization approaches.

The firm's goal is to set a pricing policy for each product that robustly maximizes the total revenue of the firm. The prices must be non-negative, and the total demand seen by the firm should be lower than its total capacity C or else the firm will pay an overbooking fee o for each unit sold above capacity. For every unit of capacity not sold, the firm will earn a salvage value of v . We require that $o \geq v$ to guarantee the concavity of the objective function. In most practical applications, the salvage value is small and the overbooking fee is large, which makes this assumption often justified. We define $w(s, \delta)$ as the terminal value of remaining inventory, which can either be an overbooking cost or a salvage value revenue. The terminal inventory is the difference between the capacity and the number of units sold: $C - \sum_{t=1}^T \tilde{d}_t(s, \delta)$. The terminal value is given by the following:

$$w(s, \delta) = -o \max\left\{\sum_{t=1}^T \tilde{d}_t(s, \delta) - C, 0\right\} + v \max\left\{C - \sum_{t=1}^T \tilde{d}_t(s, \delta), 0\right\}. \quad (1)$$

For a given pricing policy s and an uncertainty realization δ , the revenue of the firm is:

$$\Pi(s, \delta) = \sum_{t=1}^T p_t(s, \delta) \tilde{d}_t(s, \delta) + w(s, \delta). \quad (2)$$

As stated before, our goal is to find a pricing policy that will provide a robust performance for all possible demand realizations. One can think of the robust pricing problem as a game played between the firm and nature. The firm chooses a pricing policy s and nature chooses the deviations $\delta \in U$ that will

minimize the firm’s revenue. The firm seeks then to find the best robust policy. To express the different types of robust objectives explored in this study, we define $h^{obj}(s, \delta)$ as the objective function realization for given s and δ . The index *obj* can be replaced with one of three types of robust objectives: the MaxMin, the MinMax Regret and the MaxMin Ratio. For example in the MaxMin case, the objective function is simply given by $h^{\text{MaxMin}}(s, \delta) = \Pi(s, \delta)$. The remaining two objectives are explained below. The following model defines the robust pricing model:

$$\begin{aligned} & \max_{s \in S, z} z \\ \text{s.t. } & \left\{ \begin{array}{l} z \leq h^{obj}(s, \delta) \\ p(s, \delta) \geq 0 \end{array} \right\} \forall \delta \in U. \end{aligned} \tag{3}$$

As a drawback, the MaxMin approach often finds conservative pricing policies. To avoid this issue, we also consider the MinMax Regret and MaxMin Ratio robust approaches. In these cases, the firm wants to minimize the regret it will incur from using a certain policy relative to the best possible revenue in hindsight, that is, after observing the demand realization. In other words, we define the optimal hindsight revenue $\Pi^*(\delta)$ as the optimal revenue the firm could achieve, if it knew the demand uncertainty beforehand:

$$\Pi^*(\delta) = \max_{y \geq 0} \sum_{t=1}^T y_t \tilde{d}_t(y, \delta) + w(y, \delta). \tag{4}$$

The model above for $\Pi^*(\delta)$ is a deterministic convex optimization problem, which can be computed efficiently for any given δ . We next define the absolute regret as the difference between the hindsight revenue and the actual revenue, that is, $\Pi^*(\delta) - \Pi(s, \delta)$.¹ Finally, the MaxMin Ratio seeks to bind the ratio between the actual and hindsight revenues, that is, $h^{\text{Ratio}}(s, \delta) = \frac{\Pi(s, \delta)}{\Pi^*(\delta)}$. The inputs to the general robust model are as follows: the structure of the pricing policy $p_t(s, \delta)$, the demand functions $\tilde{d}_t(s, \delta)$, the objective function $h^{obj}(s, \delta)$, and the uncertainty set U (which we will replace with samples from U in section 2.2). The outputs are the set of pricing decisions s and the variable z (the robust objective value).

Note that when using adjustable policies, the profit function $\Pi(s, \delta)$ is still convex in s for any fixed realization of δ . On the other hand, the function $\Pi(s, \delta)$ is neither concave nor convex with respect to δ . In Appendix B, we illustrate this fact with a simple instance of the adjustable MaxMin model. Because of the lack of convexity or concavity, the traditional robust optimization methods (i.e., solving the robust problem using duality arguments

or simply searching over the boundary of the uncertainty set) are intractable. In the next section, we introduce the sampling-based approach for solving the robust pricing problem.

2.2. Sampling-Based Optimization

Ideally, we would like to solve the exact robust pricing problem (3), but as we mentioned, this can be intractable. Instead, assume that we are given N possible uncertainty scenarios $\delta^{(1)}, \dots, \delta^{(N)}$ where each realization $\delta^{(i)}$ is a vector containing a value for each time period and product. We use the given sampled scenarios to approximate the uncertainty set, replacing the continuum of constraints in Equation (3) by a finite number of constraints. It is natural to question how good is the solution of the sampling-based problem relative to the original problem. We next present the formulation of the sampling-based optimization problem:

$$\begin{aligned} & \max_{s \in S, z} z \\ \text{s.t. } & \left\{ \begin{array}{l} z \leq h^{obj}(s, \delta^{(i)}) \\ p(s, \delta^{(i)}) \geq 0 \end{array} \right\} \forall i = 1, \dots, N. \end{aligned} \tag{5}$$

The idea of using sampled uncertainty scenarios, or data points, in stochastic optimization models is often called sampling (or scenario)-based optimization. For tractability purposes, we impose the following linearity assumption on the set of pricing policies.

ASSUMPTION 2. *We restrict our pricing function $p_t(s, \delta_1, \dots, \delta_{t-1})$ to policies that are linear/affine on s for any given δ . We also assume that s is restricted by the strategy space S , which is a finite dimensional, compact and convex set.*

Note that in general, we still allow the dependence on δ to be nonlinear. For example, the pricing function used in our numerical studies presented in section 3 is given by:

$$p_t(s_{1,t}, s_{2,t}, \delta_1, \dots, \delta_{t-1}) = s_{1,t} + s_{2,t} \sum_{t'=1}^{t-1} \delta_{t'}. \tag{6}$$

In words, $s_{1,t}$ is the static component of price (i.e., the intercept), and $s_{2,t}$ adjusts the price linearly according to the cumulative deviation from the expected sales at time t . As long as the dependence on s is linear, this function could be refined, using for example δ_t^2 . Note that under Assumptions 1 and 2, the constraints in (5) define a convex feasible set for any given vector δ . As a result, the scenario-based problem is a convex optimization problem

and can be solved by any nonlinear optimization solver. In the next section, we present simulation results in order to validate and test the performance of our models.

3. Model Performance

The optimization problems discussed in this section were solved using the AMPL modeling language and the LOQO solver. The following experiments are designed to compare the out-of-sample performance of the different objectives discussed in section 2, and to draw some insights into the resulting pricing policies. More specifically, we compare the following models: MaxMin, MinMax Regret, MaxMin Ratio, SAA and a Bayesian approach. Note that we introduce the SAA and the Bayesian approach only as benchmarks to our robust models, and we will explain them in detail later in this section. The revenue performance of these models will be measured both in average terms and in robustness. We formally measure robustness as the conditional value-at-risk (CVaR), further explained in section 3.1.

In order to isolate the differences between the objectives, we consider an instance with a linear demand function and an additive uncertainty. In particular, the set of parameters used in this section are: $T = 2$, $d_t(p_t) = 200 - 0.5p_t$, $C = 120$, $o = 1000$, $v = 10$. For simplicity, we consider a box-type uncertainty set, $U = \{\delta: |\delta_t| \leq 15, \forall t = 1, 2\}$. In section 3.1, we consider a uniform uncertainty distributed over the entire set and we sample the data points directly from this distribution. In section 3.2, we still use the uniform distribution for sampling and optimization purposes, but define the true underlying distribution to be a truncated normal distribution. This Random Scenario Sampling approach (i.e., using an artificial sampling distribution), will be further explained in section 3.2 and theoretical guarantees are developed in section 4. We use the adjustable pricing policy displayed in Equation (6). An extension with a nonlinear demand is considered in section 3.3. Finally, section 3.4 compares our robust models to a Bayesian approach. We have performed a variety of experiments, with different demand functions, sampling distributions and pricing functions. The qualitative insights with these alternative experiments were the same, and therefore we will report only some of them in this study.

As discussed, the MaxMin model maximizes the worst-case revenue $\Pi(s, \delta^{(i)})$ for each data point $\delta^{(i)}$. For the regret-based models, we first compute the hindsight revenue $\Pi^*(\delta^{(i)})$ for each data point $\delta^{(i)}$ by solving a convex optimization problem. The MinMax Regret model maximizes the worst-case $\Pi(s, \delta^{(i)}) - \Pi^*(\delta^{(i)})$, whereas the MaxMin Ratio maximizes the worst-case $\Pi(s, \delta^{(i)})/\Pi^*(\delta^{(i)})$. As we

mentioned, robust optimization is often criticized for generating conservative policies. In other words, it provides good performance in the worst-case, but bad overall average performance. For this reason, we will compare the solution of our robust models to a more traditional stochastic optimization framework, where we maximize the expected revenue (SAA). In this case, the objective approximates the expected revenue by taking the average of each point: $h^{SAA}(s, \delta^i) = \frac{1}{N} \sum_{i=1}^N (\sum_{t=1}^T p_t(s, \delta^i) d_t(s, \delta^i) + w(s, \delta^i))$.

Note that the intrinsic volatility of the demand creates too much noise to make precise qualitative judgments on all the models simultaneously. This issue can be fixed by performing pairwise tests between the different models, which compare the difference in revenues evaluated at each data sample. The difference operator allows to cancel out the volatility in revenues, and hence yields a better comparison of the models' performances.

3.1. Pairwise Comparisons

We next perform pairwise comparisons between the MaxMin Ratio and the other models. We decided to focus on the MaxMin Ratio model as it demonstrated to be the most balanced between the robust models and the SAA, as we will show next. The other pairwise combinations were also experimented, but are not reported as they do not provide any further insight. We measure the average difference in revenues and construct an empirical confidence interval in order to draw conclusions about which model performs better in terms of average revenue. Another useful statistic to assess robust models is the CVaR, which is also known in finance as mean shortfall.² Since we are not looking at a particular revenue distribution, but instead at the difference between two revenues, it is natural to look at both tails of the distribution. In a slight abuse of notation, we define for the positive side of the distribution, $\alpha > 50\%$, $CVaR_\alpha$ to be the expected value of a variable conditional on being above its α percentile. More formally, if X is a random variable (e.g., firm's revenue) with a cumulative distribution F , then $F^{-1}(\alpha)$ is the α percentile of the possible realizations of X . For any given α , we then define the CVaR as follows:

$$\begin{aligned} \text{For } \alpha \leq 50\%, \text{ CVaR}_\alpha(X) &= E[X|X \leq F^{-1}(\alpha)], \\ \text{For } \alpha > 50\%, \text{ CVaR}_\alpha(X) &= E[X|X \geq F^{-1}(\alpha)]. \end{aligned}$$

Another important comparison is how the different models behave with small versus large sample sizes. For this reason, we perform the experiments in the following order, starting with $N = 5$ and increasing until $N = 200$: (i) Take a sample with N uncertainty data points; (ii) Optimize each model to obtain pricing policies; (iii) Measure the out-of-sample revenues for

2000 new samples, keeping track of pairwise differences in revenues; (iv) Record statistics of the average difference, and the CVaR of the difference in revenues for each pair of models; (v) Repeat steps 1–4 for 1000 iterations to build confidence intervals for the average difference and the CVaR of the difference; (vi) Increase the size of the sample size N , and repeat steps (i)–(v).

3.1.1. Comparing SAA and MaxMin Ratio. We first compare the revenues of the SAA and the MaxMin Ratio models. In Figure 1, we display a histogram of the out-of-sample difference in revenues for a particular instance using $N = 200$ samples. More specifically, for each of the 2000 out-of-sample revenue outcomes, we calculate the difference between the SAA and the MaxMin Ratio revenues. We obtain a mean revenue difference of 133 and a standard deviation of 766. In fact, we observe that the SAA yields a small revenue advantage in many cases, but can perform much worse in a few of the bad scenarios. We also obtain $\text{CVaR}_{5\%} = -2233$, and $\text{CVaR}_{95\%} = 874$. One can see that the downside shortfall of using the SAA policy versus the MaxMin Ratio can be

damaging in the bad cases, while in the better cases the positive upside is limited.

So far, we have only demonstrated the revenue performances for a single illustrative instance of the problem. The next step is to back-up these observations by repeating the experiment multiple times and building confidence intervals. We are also interested in finding out if these differences remain when the number of samples changes. More particularly, we want to know if the average revenue difference between the SAA and the MaxMin Ratio is significantly different from zero and if the $\text{CVaR}_{5\%}$ is significantly larger than $\text{CVaR}_{95\%}$ (in absolute terms). In the first three columns of Table 1, we report the mean average difference between the SAA and MaxMin Ratio revenues, with a 90% confidence interval calculated from the percentile of the empirical distribution over 1000 iterations. One can see that for $N = 150$ and $N = 200$ the average difference is clearly positive. Consequently, *with large enough sample size, the SAA performs better on average than the MaxMin Ratio.* Observe, however, that when N decreases, the average difference moves closer to 0. When $N = 5$ and $N = 10$, the average difference in revenues becomes negative, although not statistically significant as the confidence interval contains 0. This implies that *with a small sample size, the MaxMin Ratio often performs better on average.*

In Table 1, we also report the lower 5% and upper 95% tails of the revenue difference distribution. One can see that when the sample size is small, the two sides of the distribution behave similarly. As N increases, one can see that the lower tail of the revenue difference distribution is not significantly affected, and the length of the confidence interval seems to decrease around similar mean values of $\text{CVaR}_{5\%}$. The upper tail of the distribution, on the other hand, seems to shrink as N increases. When $N = 200$, the upside benefits of using SAA versus MaxMin Ratio will be limited to an average of 691 on the 5% better cases, while the 5% worst-cases will fall short by an average of -2418 . Moreover, with a 90% confidence level, we observe that the $\text{CVaR}_{5\%}$ will be larger in absolute value than the $\text{CVaR}_{95\%}$, since the

Figure 1 Comparison between SAA and MaxMin Ratio [Color figure can be viewed at wileyonlinelibrary.com]




Table 1 Average and CVaR of the Difference: Revenue SAA—Revenue MaxMin Ratio

N	Average		5% CVaR		95% CVaR	
	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-79	[-1763, 909]	-2686	[-11,849, -198]	2469	[676, 6437]
10	-29	[-837, 549]	-2345	[-7195, -348]	2113	[609, 4115]
50	78	[-115, 324]	-2260	[-4798, -382]	962	[300, 1953]
100	103	[-17, 252]	-2355	[-4124, -735]	773	[310, 1372]
150	122	[8, 253]	-2378	[-3758, -992]	721	[344, 1174]
200	130	[31, 249]	-2418	[-3606, -1257]	691	[363, 1068]

confidence intervals do not intersect. In summary, *the SAA revenues will do much worse than the MaxMin Ratio in the low revenue cases, while not so much better in the high revenue cases.*

To summarize the pairwise comparison between SAA and MaxMin Ratio, we observe that in small sample sizes, the MaxMin Ratio seems to perform both better on average and in the worst cases of the revenue distribution, but the statistical tests cannot either confirm or deny this finding. For large sample sizes, the results are clearer: with 90% confidence level, the SAA will yield a better average revenue, but will pay a heavy penalty in the worst cases. This last conclusion can be viewed as intuitive, since the SAA model tries to maximize average revenues, while the robust model (in this case, the MaxMin Ratio) aims to protect against bad revenue scenarios. However, the better performance of the robust model over the SAA using small sample sizes is not intuitive, and understanding this behavior could lead to interesting future research.

In Appendices C.1.1 and C.1.2, we present the other pairwise comparisons between the robust models. The summary of our findings will be presented in section 3.5. So far, we have assumed that the data samples originate from the true distribution, which could be the case if we have a large pool of historical data. We next diverge from this assumption and explore the use of an artificial distribution.

3.2. Using Random Scenario Sampling

In this section, we consider the case when the sampling distribution is not the same as the true underlying distribution. More details about the theory behind this approach are provided in section 4. In contrast with the previous section, we will use a truncated normal distribution, that is, the underlying true distribution is normal with mean 5 and standard deviation 7.5 but truncated over the range $[-15, 15]$. On the other hand, we assume that the firm does not know this true underlying distribution, and instead uses a uniform distribution over the uncertainty set $[-15, 15]$. These uniformly sampled points will then be used to run our sampling-

based pricing models. Table 2 reports the pairwise comparison of the difference in revenues between the SAA and the MaxMin Ratio. *Note that the average of the difference in revenues favors the negative side, that is, the side in which the MaxMin Ratio solution outperforms the SAA. This suggests that the robust solution obtained using the MaxMin Ratio will perform better than the SAA even in terms of average revenue.* This result, although counter intuitive, appears consistent for both small and large sample sizes. The confidence intervals do not confirm this finding as statistically significant at 90% confidence level, but these intervals appear to move away from the positive side as the sample size increases. Table 3 displays the results of the pairwise comparison of the revenues for the MaxMin Ratio and the MaxMin models. One can see that in this case, the MaxMin Ratio performs better both on average and in the worst-case. The confidence interval of the average difference in revenues is positive even for the small sample size of $N = 10$, distancing away from zero as N increases. As for the CVaR comparison, note that for $N \geq 100$ the 5% better cases of the MaxMin Ratio are on average better by 2989, while the 5% better cases for the MaxMin are only better by an average of 1001. *In addition, the confidence intervals around these values do not overlap, strongly suggesting that the MaxMin Ratio is more robust than the MaxMin.* It is worth mentioning that the results of the MinMax Regret model led similar outcomes as the MaxMin Ratio, as expected given the similar nature of these two models. Between these models, there is an indication that the MaxMin Ratio performs better on average than MinMax Regret, but the MinMax Regret is more robust. To avoid redundancy, these results are not reported here.

To summarize the experiments with randomized sampling: We have used an artificial sampling distribution (uniform), while the true distribution was a biased truncated normal. We observed that the SAA usually performs worse than the MaxMin Ratio in terms of average revenue. The MaxMin Ratio model clearly outperforms the MaxMin model, both in average and worst-case differences in revenues.

Table 2 Using Artificial Sampling Distribution: Average and CVaR of the Difference Revenue SAA—Revenue MaxMin Ratio

N	Average		5% CVaR		95% CVaR	
	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-165	[-3753, 1649]	-3620	[-15,124, -318]	2997	[903, 6895]
10	-6	[-1726, 925]	-2400	[-8952, -401]	2720	[760, 5236]
50	-88	[-472, 112]	-1462	[-3748, -262]	1195	[280, 2588]
100	-31	[-192, 59]	-941	[-2848, -200]	961	[285, 2117]
150	-21	[-118, 36]	-692	[-2008, -203]	827	[177, 1849]
200	-29	[-158, 19]	-872	[-2370, -195]	670	[200, 1474]

Table 3 Using Artificial Sampling Distribution: Average and CVaR of the Difference Revenue MaxMin Ratio—Revenue MaxMin

N	Average		5% CVaR		95% CVaR	
	Mean	95% CVaR	Mean	90% CI	Mean	95% CI
5	875	[−417, 2704]	−1326	[−4438, −227]	4064	[1156, 10,306]
10	701	[22, 1601]	−1385	[−3824, −447]	3603	[2020, 6349]
50	580	[247, 1077]	−1191	[−2603, −703]	3052	[1722, 4614]
100	644	[355, 1089]	−1001	[−1553, −737]	2989	[1727, 4718]
150	672	[361, 1098]	−962	[−1194, −724]	3015	[1764, 4753]
200	659	[389, 1082]	−987	[−1338, −798]	2995	[1702, 4783]

3.3. Extensions

3.3.1. Nonlinear Demand Models. The numerical experiments we presented so far were performed using linear demand functions, although our framework can deal with nonlinear demands. In this section, we show an application of our approach with a logarithmic demand defined as $d_t(p_t) = a_t - b_t \log(p_t) + \delta_t$, that is, the demand is a linear function of the logarithm of the price plus some deviation δ_t . Using this demand function, we can apply the exact same underlying models as before. We consider $T = 2$, $a_1 = a_2 = 10$, $b_1 = b_2 = 2$, $C = 10$, and δ_t uniformly distributed over $U = \{\delta: |\delta_t| \leq 15, \forall t = 1, 2\}$. Figure 2 displays the overall performance of the four models under a logarithmic demand function. One can see that the results for the logarithmic demand model yield very similar insights and findings as the linear demand function. As we have previously explained, there is no particular difference in the methodology needed to implement different demand models, as long as they satisfy Assumption 1. If we are provided with historical data of prices and demand realizations, one should select the demand function that provides the best fit to the data.

Subsequently, the uncertainty realizations will be built from the estimation errors.

An additional extension that considers a network example with multiple products and resources is presented in Appendix C.2. We obtained that in general, the results and insights for the network example are consistent with the single product case.

3.4. Comparing to a Bayesian Approach

In this section, we consider a simple Bayesian approach as an additional benchmark. More precisely, as new samples are observed, we update the mean of the error term, δ_t , by using Bayesian updating (note that this differs from the traditional Bayesian approach which updates the prior distribution and not only the mean). In this method, we start at $t = 1$, with a given prior distribution over the uncertainty set $U = \{\delta_1: |\delta_1| \leq 15\}$. We then update the mean of the prior distribution as more data samples become available at the end of the first period. In each period t , after observing the N samples, we compute the sample average of the N data points: $\mu_t = \frac{\sum_{i=1}^N \delta_t^{(i)}}{N}$. We then update the mean of the posterior distribution as follows:

$$E[\delta_{t+1} | \delta_t] = \theta E[\delta_t] + (1 - \theta) \mu_t, \quad (7)$$

where $0 \leq \theta \leq 1$ is a given parameter. When θ is close to 1, it means that we give a high weight (or belief) to the prior distribution, whereas when θ is close to 0, we give a high weight to the realized samples from the previous period. We then solve the optimization problem with the expected demand in the objective function. More precisely, for each time period, we compute the expectation by using the updated distribution in Equation (7). Note that we impose the constraints in the optimization problem to remain the same as before. We do so for two reasons: (i) We want to enforce the capacity constraints using the realized demand rather than the expected values; and (ii) This allows us to compare the outcomes of the Bayesian and robust approaches since they both have the same feasible region.

Figure 2 Logarithmic Demand Model: $T = 2$, $C = 10$, $a_t = 10$, $b_t = 2$
[Color figure can be viewed at wileyonlinelibrary.com]




Table 4 Average and CVaR of the Difference: Revenue MaxMin—Revenue Bayesian with $E[\delta_1] = 0$ and $\theta = 0.5$

N	Average		95% CVaR		90% CI	
	Mean	90% CI	Mean	90% CI	Mean	95% CVaR
5	358	[−3807, 4672]	−8221	[−10,951, −6389]	9609	[7666, 11,887]
10	1628	[−2161, 8389]	−6496	[−9476, −4164]	12,847	[10,904, 15,348]
20	1961	[−2018, 8984]	−5552	[−8787, −3350]	13,947	[11,747, 17,339]
150	2316	[−1895, 9046]	−3283	[−4258, −2576]	14,132	[11,684, 17,318]

Table 5 Average and CVaR of the Difference: Revenue MaxMin Ratio—Revenue Bayesian with $E[\delta_1] = 0$ and $\theta = 0.5$

N	Average		5% CVaR		95% CVaR	
	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	712	[−3012, 5471]	−7642	[−10,938, −5546]	11,161	[9012, 14,393]
10	1733	[−1775, 8191]	−5797	[−8860, −3264]	11,543	[10,245, 13,119]
20	2047	[−1799, 8949]	−4793	[−8263, −2632]	13,025	[11,208, 15,790]
150	2387	[−1773, 9611]	−2665	[−3276, −2239]	14,929	[12,481, 18,151]

In order to implement this Bayesian approach, we need to know the mean of the prior distribution, denoted by $E[\delta_1]$ and the updating parameter θ . We test three different prior distributions: $E[\delta_1] = -5, 0, 5$ as well as three different values for the parameter $\theta = 0.2, 0.5, 0.8$. The results when we use the true prior (i.e., $E[\delta_1] = 0$) and $\theta = 0.5$ for the MaxMin and MaxMin Ratio objectives are presented in Tables 4 and 5, respectively. The other results can be found in Appendix C.3.

One can see that the two robust models outperform the Bayesian approach both in terms of average revenue and worst-case values. The average of the mean is always positive with a 90% confidence interval skewed toward the positive side (it still includes 0 though). Furthermore, the 5% best cases for any of the two robust models (95% CVaR column) are significantly better than the 5% best cases of the Bayesian approach (5% CVaR column). Note that for $N \geq 10$, the confidence intervals for the 95% and 5% CVaR do not intersect, suggesting that with a 90% confidence level, the MaxMin and MaxMin Ratio models are more robust than the Bayesian approach. In addition, when we use a wrong value of the prior (i.e., $E[\delta_1] \neq 0$), the superiority of the robust models can be even more significant (as we can see from the tables presented in Appendix C.3). These tests suggest that the robust models proposed in this study outperform the Bayesian approach we considered, even when the true prior distribution is known.

3.5. Summary of Results from Numerical Experiments

The key insights of our numerical experiments can be summarized as follows.

- The SAA has a risk-neutral perspective, thus maximizing average revenues, regardless of the final revenue distribution. It yields the best average revenue performance if the number of samples is large enough, and if they are reliably coming from the true distribution. Nevertheless, the revenue outcome is subject to the largest amount of variation, with a large shortfall in the bad cases when compared to the robust models.
- The MinMax Regret model seems to attain the most robust revenue performance when compared to the other models. The MaxMin Ratio strikes a balance between the conservativeness of the MinMax Regret, and the aggressiveness of the SAA.
- All the robust models (MaxMin, MinMax Regret and MaxMin Ratio) perform just as well as the SAA in average revenues for small sample sizes. For large sample sizes, the robust models have smaller average revenues, but also smaller shortfalls for the extreme revenue cases, that is, good for risk-averse firms. Note that the advantage of having more stable revenue distributions are present even for small sample sizes.
- The MaxMin model tends to be too conservative. It is usually dominated by the MaxMin Ratio model both in terms of worst-case and average revenues.
- The MaxMin Ratio and MinMax Regret have the most consistent performance when the data set does not come from the true underlying distribution.
- Both the MaxMin and the MaxMin Ratio models seem to outperform a simple Bayesian approach based on iteratively updating the mean of the demand uncertainty.

4. Theoretical Guarantees

Our goal is to relate the sampled problem (5) to the original robust problem (3). To this end, we first need to formally show the existence of an optimal solution for the original robust problem. Next, we demonstrate that the sampled problem solution converges to the optimal solution almost surely, as the number of samples goes to infinity. We then derive a bound on the sampling size, and provide probabilistic guarantees, adapted from Calafiore and Campi (2006), showing that the sampled solution is ε -robust. Finally, we discuss some interesting properties of the dual sampled problem. By exploiting the structure of the optimal solution of the dual sampled problem, we can identify the critical demand samples that ultimately drive the optimal policy.

4.1. Existence and Convergence

In order to later show convergence of the sampled problem to the original robust problem, we first verify that the original robust problem in (3) has an optimal solution.

PROPOSITION 1. *There exists an optimal solution to Equation (3).*

This result relies on the continuity of the objective. For more details, see Appendix D. We next show that the sampled problem converges to that solution when N goes to infinity. In order to simplify notation, we define the decision variable $x = (s, z)$ as the combined pricing policy decision s and objective value z . Note that it lies within the domain X , that is, $x \in X$, where $X = S \times \mathfrak{R}$. Define c with the same dimension as x such that $c = (0, \dots, 0, 1)$. Using x and c , we can write the pricing problem using a standard convex optimization form with a linear objective. Define the equivalent constraint function with a scalar valued g given by: $g(x, \delta) = \max\{z - h^{obj}(s, \delta), -p(s, \delta)\}$. Since each constraint in the definition of $g(x, \delta)$ is convex in x for a given δ , the maximum is also convex in x . Moreover, $g(x, \delta) \leq 0$ is equivalent to the constraint set in (r). Then, problem (r) can be written as:

$$\max_{x \in X} c'x, \quad \text{s.t. } g(x, \delta) \leq 0, \quad \forall \delta \in U. \quad (8)$$

Since we cannot solve the problem with a continuum of constraints, we instead consider the problem with a finite number of deviations $\delta^{(i)}$ from the uncertainty set U , where $i = 1, \dots, N$. Then, the concise version of the sampled problem (r) can be written as:

$$\max_{x \in X} c'x, \quad \text{s.t. } g(x, \delta^{(i)}) \leq 0, \quad \forall i = 1, \dots, N. \quad (9)$$

The following assumptions are needed to show the convergence of the sampled problem solution to the

optimal robust solution, when the number of samples goes to infinity. Denote the feasible region of the robust problem (r) as \bar{X} , that is, $\bar{X} = \{x \in X : g(x, \delta) \leq 0, \delta \in U\}$. Similarly, let $\bar{X}^N = \{x \in X : g(x, \delta^{(i)}) \leq 0, i = 1, \dots, N\}$ be the feasible region of (r).

ASSUMPTION 3.

- (i) *The feasible region \bar{X} of Equation (8) is convex and compact.*
- (ii) *The samples $\delta^{(i)}$ are independently drawn from a compact uncertainty set U .*
- (iii) *The PDF of the sampling distribution, denoted by f , is bounded below, that is, there exists $\varepsilon > 0$ such that $f(\delta) > \varepsilon$ for all $\delta \in U$.*

We denote the optimal objective value of Equations (8) and (9) by z^* and z^N , respectively. Note that for each N , z^N is a random variable since $\delta^{(i)}$ are drawn from the uncertainty set U .

THEOREM 1. *Under Assumption 3, the sampled optimal solution x^N and optimal value z^N converge almost surely to the optimal solution x^* and optimal value z^* , respectively.*

We note that the result of Theorem 1 might appear intuitive but a formal proof is not straightforward. It implies that a series of convex problems converge to a non-convex problem with probability 1. This convergence result is the first natural validation that the sampled problem can approximate its robust counterpart.

4.2. Probabilistic Guarantees of the Sampled Solution

For any practical application, it is important to understand what will happen to the sampled solution when only a finite number of samples is available. Depending on what type of demand information and historical data provided to the firm, we propose two solution approaches: the Data-driven and the Random Scenario Sampling. The first approach, Data-driven, assumes that we are given a large pool of uncertainty data drawn from the true underlying distribution, for example, from historical data. The second approach, Random Scenario Sampling, assumes that we do not have enough (if any) data points from the true distribution, but instead we have a sampling distribution which can be used to generate random data points. In either case, the following definition of ε -robustness, from Calafiore and Campi (2005), is required to develop a sampling size bound.

DEFINITION 1. *For a given pricing policy x and a distribution \mathcal{Q} of the uncertainty δ , we define the probability of violation $V_{\mathcal{Q}}(x)$ as: $V_{\mathcal{Q}}(x) = \mathbf{P}_{\mathcal{Q}}\{\delta : g(x, \delta) > 0\}$.*

Note that the probability of violation corresponds to a measure on the actual uncertainty realization δ , which has an underlying unknown distribution \mathcal{Q} . In other words, for the MaxMin objective, given the pricing policy $x = (s, z)$, $V_{\mathcal{Q}}(x)$ is the probability that the actual demand realization yields a revenue lower than z , that is, the worst-case revenue. Therefore, it is easy to interpret any violation as an unexpected loss in revenue. An ϵ -robust feasible solution is then defined as follows.

DEFINITION 2. *If $V_{\mathcal{Q}}(x) \leq \epsilon$, then x is ϵ -level robustly feasible (or simply ϵ -robust).*

Note that the given set of scenario samples is random and comes from the probability space of all the possible sampling outcomes of size N . For a given $\epsilon \in (0, 1)$, a “good” sample is such that the solution $x^N = (s^N, z^N)$ of the sampled problem will lead to an ϵ -robust solution, that is, the probability of nature giving the firm some revenue below our estimated z^N is smaller than ϵ . Define the confidence level $(1 - \beta)$ as the probability of sampling a “good” set of scenario samples. Alternatively, β is known as the “risk of failure,” which is the probability of drawing a “bad” sample. Our goal is to determine the relationship between the confidence level $(1 - \beta)$, the robust level ϵ , and the number of samples N .

Before describing the sampling size bound, there is one last concept that is required for the Random Scenario Sampling approach. Suppose that we do not have samples obtained from the true distribution \mathcal{Q} . Instead, we are given the nominal demand parameters and the uncertainty set U . We would like to draw some samples by using another distribution \mathcal{P} and run the sampling-based pricing model in (9). In order to make a statement about the confidence level of the solution and the sample size, we impose an assumption about how close \mathcal{P} is to the true distribution \mathcal{Q} .

DEFINITION 3. (BOUNDED LIKELIHOOD RATIO). *We say that the distribution \mathcal{Q} is bounded by \mathcal{P} with factor k if for every subset A of the sample space: $\mathbf{P}_{\mathcal{Q}}(A) \leq k\mathbf{P}_{\mathcal{P}}(A)$.*

In other words, the true unknown distribution \mathcal{Q} does not have concentrations of mass that are unpredicted by \mathcal{P} . If $k = 1$, then the two distributions are the same, except for a set of probability 0, and therefore the scenario samples come from the true distribution. Note that the definition above is satisfied under a more restrictive, but perhaps more common assumption for continuous distributions of Bounded Likelihood Ratio $\frac{d\mathbf{P}_{\mathcal{Q}}(x)}{d\mathbf{P}_{\mathcal{P}}(x)} \leq k$.

On one hand, it seems hard for a manager to select a bound k on the likelihood ratio that would work for the uncertainty set. On the other hand, the variance of

the demand is a popular measure to most managers and might be somehow estimated by the firm. Gaur et al. (2007) proposed a way to estimate the variance of the demand distribution by using the dispersion of managers’ judgmental forecasts. With the variance of the demand distribution and the additional assumption that the demand uncertainty is independent across time periods, Lobel (2012) proposed a way to derive a likelihood ratio bound k between a uniform distribution and any log-concave distribution. Other measures for k can be obtained by bounding the mode of the demand distribution. For the remainder of the study, we assume that such a bound k is known to the firm.

Let n_x be the dimension of the strategy space X . Theorem 2 adapts the sampling size bound from Calafiore and Campi (2006) to allow sampling from an artificial distribution. It provides a bound on the probability that the solution of the sampled problem is not ϵ -robust, that is, the probability of drawing a “bad” sample.

THEOREM 2. *Assume the sampling distribution \mathcal{P} bounds the true distribution \mathcal{Q} by a factor of k (see Definition 3). The “risk of failure” parameter $\beta(N, \epsilon)$ can be defined as:*

$$\beta(N, \epsilon) \doteq \binom{N}{n_x} (1 - \epsilon/k)^{N-n_x}.$$

Then, with probability greater than $[1 - \beta(N, \epsilon)]$, x^N is ϵ -level robustly feasible, that is,

$$\mathbf{P}_{\mathcal{P}}((\delta^{(1)}, \dots, \delta^{(N)}) : V_{\mathcal{Q}}(x^N) \leq \epsilon) \geq 1 - \beta(N, \epsilon).$$

In other words, the level $\beta(N, \epsilon)$ is a bound on the probability of getting a “bad” sample of size N for a given robust level ϵ . Therefore, $1 - \beta(N, \epsilon)$ is the confidence level that our solution is ϵ -robust. As a corollary of Theorem 2, one can obtain a direct sample size bound for a desired confidence level β and robust level ϵ .

COROLLARY 1. *Adapted from Calafiore and Campi (2006), if the sample size N satisfies: $N \geq N(\epsilon, \beta) \doteq (2k/\epsilon) \ln(1/\beta) + 2n_x + (2n_x k/\epsilon) \ln(2k/\epsilon)$, then, with probability greater than $1 - \beta$ the solution of the sampled problem is ϵ -robust.*

Note that the bound in Corollary 1 is not necessarily the tightest value of N that satisfies Theorem 2 for given β and ϵ . Numerically solving for N the equation $\beta = \binom{N}{n_x} (1 - \epsilon/k)^{N-n_x}$ may yield a smaller sample size requirement. Nevertheless, Corollary 1 offers a direct calculation and insight into the relationship between the sample size and the confidence level of the sampled solution. Note that $N(\epsilon, \beta)$ goes to

infinity as either ϵ or β goes to zero, which is rather intuitive. Note also that the dependence on β is of the form $\ln(1/\beta)$, meaning that the confidence parameter β can be pushed down toward zero without significantly affecting the number of samples required. For implementation purposes, it allows us to keep a good level of confidence and design N based on the desired level of robustness.

The difference between Theorem 2 above and Theorem 1 from Calafiore and Campi (2006) is the introduction of randomized sampling. In order to apply the above confidence bound to our pricing problem, we only need to know k and n_x , that is, the number of decision variables in the optimization model. The latter depends on the functional form of the pricing policy under consideration. For example, in the case of static pricing (open-loop), the number of decision variables is the number of time periods T plus one for the robust objective variable, that is, $n_x = T + 1$. If the policy is adjustable, it requires additional decision variables. In particular, for a price that is a linear function of the sum of previously realized uncertainties, as in section 3, $n_x = 2T + 1$.

Note that this Random Scenario Sampling result can be quite useful in practice, especially in settings where there is a small amount of data or no data at all. To the best of our knowledge, this is the first result that gives a robust sampling size bound when using an artificial sampling procedure and such limited knowledge of the true distribution.

4.3. Dual Sampled Problem

Consider the sampled problem in (5). We introduce the dual variables $\lambda_1, \lambda_2, \dots, \lambda_N \geq 0$ for each constraint $z \leq h^{obj}(s, \delta^{(i)})$, $i = 1, 2, \dots, N$. Standard arguments from Lagrange duality imply that the dual objective is an upper bound to the optimal objective value of problem (5). Furthermore, since the primal variable z is unconstrained, we require that $\sum_{i=1}^N \lambda_i = 1$ for this upper bound to be finite. Consequently, the dual variables $\lambda_1, \lambda_2, \dots, \lambda_N$ represent probability distributions over sampled residuals. In addition, since the sampled problem (5) is concave with respect to the variables (s, z) , it implies that strong duality holds. We denote the optimal dual solution by $\lambda_1^*, \lambda_2^*, \dots, \lambda_N^*$, so that the optimal primal objective can be expressed as the following convex combination: $z^* = \sum_{i=1}^N \lambda_i^* h^{obj}(s^*, \delta^{(i)})$.

The above duality argument shows that in our sampled problem, nature chooses a probability distribution over the set of potential residuals $\delta^{(i)}$, $i = 1, \dots, N$. This probability distribution is fully characterized by the optimal dual variables. Then, the optimal pricing policy s can be obtained by optimizing against this worst-case distribution. This bears an interesting interpretation that can help the firm

understand the critical demand samples that ultimately drive the optimal pricing policy. Note that similar ideas and arguments were used in previous works (see, e.g., Lim et al. (2012) in the context of robust portfolio choice problems, and the book by Luenberger (1969)). A similar analysis can be carried out for the original robust problem (3), but requires an infinite dimensional dual space, and leads in turn to nature optimizing over probability distributions on the uncertainty set U . In addition, since the original problem in (3) is not concave (as we have shown in Appendix B), strong duality does not necessarily hold.

With the above duality argument in mind, one can use the optimal dual variables of the sampled problem to draw additional insights. For example, we are interested in studying how the different robust objectives (MaxMin, MaxMin Ratio and MinMax Regret) will affect the optimal dual solution. In other words, we want to see if the different objectives assign similar weights to the same demand realizations. Furthermore, we want to understand how the number of samples N affects the optimal dual solution. We present some computational experiments in Appendix C.4. Interestingly, we show that in most cases, the type of robust objective does not affect significantly the dual variables, and hence the critical demand realizations. That is, for a given value of N , under any of the three robust objectives, we obtain a similar probability distribution chosen by nature. More precisely, the demand realizations with a significant weight (captured by the optimal dual variables λ_i^*) do not vary much under the three different robust objectives. Note that these critical demand realizations will ultimately drive the optimal pricing policy. The relative weights may still depend on the objective type (see more details in Appendix C.4). This can be used in order to understand how previous demand realizations are affecting future pricing decisions. For example, some popular days with high demand volumes can be directly identified from the optimal dual variables. On the other hand, as we increase the number of sampled data N , we could not observe a clear pattern of convergence. Even in cases where some demand realizations are critical for a small value of N , they will not necessarily remain critical as N increases. This confirms the fact that we do not necessarily have the convergence of the dual variables, even though the primal optimal solution converges almost surely, as we have shown in Theorem 1.

5. Conclusion

In this study, we developed a framework to solve the robust dynamic pricing problem using a

sampling-based approach. Our model is a practical decision tool that can easily be implemented as a convex optimization model. Using a sampling approach allows for different objective functions and adaptable pricing strategies, hence tackling two standard issues in robust optimization: conservative solutions and open-loop policies.

We show that the regret-based objective functions deliver a more robust revenue performance relative to the sample-average approximation and to a Bayesian approach. The sample-average approximation has a better average revenue performance when a large sample of demand data is available and is obtained from the true underlying distribution. On the other hand, when the amount of data is small or sampled from an artificial sampling distribution, the regret-based models tend to perform better for the bad cases as well as in terms of average performance. This small sample size performance could make this framework a useful tool for managers with constraints on data availability.

On the theoretical side, we show that the sampled solution converges almost surely to the robust solution, and we provide performance guarantees based on the sample size. More specifically, we extend the known sampling size bound for data-driven optimization to a random scenario sampling approach. This approach connects the data-driven and the robust optimization fields. Starting with an intractable robust model, we approximate the solution with data-driven techniques, but using limited actual demand data.

Acknowledgments

The authors thank the department editor (William L. Cooper), the associate editor, and the two anonymous referees for their insightful comments, which have helped improve this study. The authors also thank Yuan Zhong for his insightful suggestions regarding the proof of Theorem 1, and Rachel Burns for her help running the computations presented in section 3.4 and in Appendices C.3 and C.4. The third author acknowledges the NSF grants CMMI-1563343 and CMMI-1162034 that made this research possible.

Notes

¹To be consistent with the MaxMin formulation, we define the objective function for the MinMax Regret as $I^{\text{Regret}}(s, \delta) = \Pi(s, \delta) - \Pi^*(\delta)$, which is the negative of the regret.

²Taken at a given percentile α , CVaR_α is the expected value of a random variable conditional that the variable is below its α percentile. This statistic is a way to measure the tail of a distribution, and allows to determine how bad things can go when they do go wrong.

References

- Adida, E., G. Perakis. 2006. A robust optimization approach to dynamic pricing and inventory control with no backorders. *Math. Program.* **107**(1): 97–129.
- Araman, V. F., R. Caldentey. 2009. Dynamic pricing for non-perishable products with demand learning. *Oper. Res.* **57**(5): 1169–1188.
- Aviv, Y., A. Pazgal. 2002. Pricing of short life-cycle products through active learning. Working paper.
- Ball, M. O., M. Queyranne. 2009. Toward robust revenue management: Competitive analysis of online booking. *Oper. Res.* **57**(4): 950–963.
- Ben-Tal, A., A. Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Math. Program.* **88**: 411–424.
- Ben-Tal, A., A. Goryashko, E. Guslitzer, A. Nemirovski. 2004. Adjustable robust solutions of uncertain linear programs. *Math. Program.* **99**(2): 351–376.
- Ben-Tal, A., B. Golany, A. Nemirovski, J. Vial. 2005. Retailer-supplier exible commitments contracts: A robust optimization approach. *Manuf. Serv. Oper. Manag.* **7**(3): 248–271.
- Bertsimas, D., M. Sim. 2004. The price of robustness. *Oper. Res.* **52**(1): 35–53.
- Bertsimas, D., A. Thiele. 2006. A robust optimization approach to supply chain management. *Oper. Res.* **54**(1): 150–168.
- Bertsimas, D., P. Vayanos. 2017. Data-driven learning in dynamic pricing using adaptive optimization. Working paper.
- Bertsimas, D., D. Iancu, P. Parrilo. 2010. Optimality of affine policies in multi-stage robust optimization. *Math. Oper. Res.* **35**(2): 363–394.
- Bertsimas, D., D. A. Iancu, P. A. Parrilo. 2011. A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Trans. Autom. Control* **56**(12): 2809–2824.
- Calafiore, G. 2009. On the expected probability of constraint violation in sampled convex programs. *J. Optim. Theory Appl.* **143**(2): 405–412.
- Calafiore, G., M. Campi. 2005. Uncertain convex programs: Randomized solutions and confidence levels. *Math. Program.* **102**: 25–46.
- Calafiore, G., M. Campi. 2006. The scenario approach to robust control design. *IEEE Trans. Autom. Control* **51**(5): 742–753.
- Campi M., S. Garatti. 2008. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM J. Optim.* **19**(3): 1211–1230.
- Caramanis, C. 2006. Adaptable optimization: Theory and algorithms. Ph.D dissertation, Department of Electrical Engineering and Computer Science, MIT.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing: Research overview, current practices and future directions. *Management Sci.* **49**: 1287–1309.
- Eren, S. S., C. Maglaras. 2010. Monopoly pricing with limited demand information. *J. Revenue Pricing Manage.* **9**: 23–48.
- Farias, V. F., B. Van Roy. 2010. Dynamic pricing with a prior on market response. *Oper. Res.* **58**(1): 16–29.
- Gallego, G., I. Moon. 1993. The distribution free newsboy problem: Review and extensions. *J. Oper. Res. Soc.* **44**(8): 5.
- Gaur, V., S. Kesavan, A. Raman, M. Fisher. 2007. Estimating demand uncertainty using judgmental forecasts. *Manuf. Serv. Oper. Manag.* **9**(4): 480–491.
- Kleywegt, A., A. Shapiro, T. Homem-de Mello. 2001. The sample average approximation method for stochastic discrete optimization. *SIAM J. Optim.* **12**(2): 479–502.

- Lim, A., J. Shanthikumar, T. Watewai. 2009. Robust multi-product pricing and a new class of risk-sensitive control problems. Working paper.
- Lim A. E., J. G. Shanthikumar, G. Y. Vahn. 2012. Robust portfolio choice with learning in the framework of regret: Single-period case. *Management Sci.* 58(9): 1732–1746.
- Lin, K. Y. 2006. Dynamic pricing with real-time demand learning. *Eur. J. Oper. Res.* 174(1): 522–538.
- Lobel, R. 2012. Pricing and incentive design in applications of green technology subsidies and revenue management. Ph.D. dissertation, Operations Research Center, MIT.
- Luenberger, D. G. 1969. *Optimization by Vector Space Methods*. John Wiley & Sons, New York.
- Pagnoncelli, B., S. Ahmed, A. Shapiro. 2009. Computational study of a chance-constrained portfolio problem. *J. Optim. Theory Appl.* 142(2): 399–416.
- Perakis, G., G. Roels. 2010. Robust controls for network revenue management. *Manuf. Serv. Oper. Manag.* 12(1): 56–76.
- Rusmevichientong P., B. Van Roy, P. W. Glynn. 2006. A non-parametric approach to multi-product pricing. *Oper. Res.* 54(1): 89–98.
- Talluri K. T., G. J. Van Ryzin. 2006. *The Theory and Practice of Revenue Management*, vol. 68. Springer Science & Business Media, New York, NY.
- Vayanos P., D. Kuhn, B. Rustem. 2012. A constraint sampling approach for multi-stage robust optimization. *Automatica* 48(3): 459–471, ISSN 0005-1098.

Supporting Information

Additional supporting information may be found online in the supporting information tab for this article:

Appendix A: Notation.

Appendix B: Illustration of the Non-Convexity Issue.

Appendix C: Additional Experiments.

Appendix D: Proof of Proposition 1.

Appendix E: Proof of Theorem 1.

Appendix F: Proof of Theorem 2.