

MIT Open Access Articles

Dynamic Pricing through Data Sampling

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Cohen, Maxime C., Ruben Lobel, and Georgia Perakis. "Dynamic Pricing through Data Sampling." *Production and Operations Management* 27, no. 6 (March 1, 2018): 1074–1088.

As Published: <http://dx.doi.org/10.1111/POMS.12854>

Publisher: Wiley

Persistent URL: <http://hdl.handle.net/1721.1/120722>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Dynamic Pricing Through Data Sampling

Maxime C. Cohen^a, Ruben Lobel^b, Georgia Perakis^c

^a*Operations Research Center, MIT, Cambridge, MA 02139, USA*

^b*The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA*

^c*Sloan School of Management, MIT, Cambridge, MA 02139, USA*

Abstract

In this paper 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 robust and dynamic pricing strategy that maximizes revenue over the entire selling season. We develop a tractable optimization model that directly uses demand data, therefore creating a practical decision tool. Furthermore, we provide theoretical performance guarantees for this sampling-based solution, based on the number of samples used. Finally, we compare the revenue performance of our model using numerical simulations, exploring the behavior of the model with different robust objectives, sample sizes, and sampling distributions. This modeling approach could be particularly important for risk-averse managers with limited access to historical data or information about the demand distribution.

Keywords: dynamic pricing, data-driven, sampling based optimization, robust optimization

1. Introduction

In many industries, managers are faced with the challenge of selling a fixed amount of inventory within a specific time horizon. Examples include the case of airlines selling flight tickets, hotels trying to book rooms and retail stores selling products for the current season. All these cases share a common trait: a fixed initial inventory that cannot be replenished within the selling horizon. The firm's goal is to set prices at each stage of the given selling horizon that will maximize revenues while facing an uncertain demand.

As in most real-life applications, the intrinsic randomness of the firm's demand is an important factor that must be taken into account. Decisions based on deterministic forecasts of demand can expose the firm to severe revenue losses. When introducing demand uncertainty in the optimization model, a typical approach is stochastic optimization. A disadvantage of this approach is that it requires full knowledge of the demand distribution. In practice, this approach can sometimes lead to very large data requirements. It also assumes the uncertainty to be independent across time periods, otherwise creating dimensionality problems.

In many application settings, firms do not have a lot of information about this demand uncertainty. In some cases, they have some historical data of price and demand levels. In other cases, the information available is perhaps limited to the range of the demand. With this in mind, our model does not make assumptions about the demand distribution. Instead, we directly use the available demand data in the optimization model, therefore creating a very practical analytic tool for pricing problems.

This distribution-free modeling approach to demand uncertainty is inspired by the recent literature in robust optimization. In robust pricing models, one does not assume a known distribution for the demand uncertainty, but assumes only that it lies within 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 about the distribution or correlation of demand 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 policy). The dynamic models that search for closed-loop policies, i.e. which account for previously realized uncertainties and adjust the pricing decisions, can easily become intractable. Our goal in this paper is to come up with an approach that can tackle these two issues, by developing a framework for finding non-conservative and adjustable robust pricing policies.

1.1. Contributions

The main contributions of this paper are:

- *It proposes a pricing decision tool that is flexible and easy to implement.*

The resulting optimization model is a tractable convex optimization problem. We accomplish this by using a sampling based optimization approach. This allows for a wide range of modeling complexity such as adjustable pricing policies, large class of nonlinear demand functions, overbooking and salvage value.

- *It directly uses available data in the price optimization model.*

With the proposed sampling based technique, we can define the pricing optimization model without imposing a probability distribution for the demand. Consequently, we admit for any arbitrary demand distribution that could potentially be correlated across time periods. In contrast to our methodology, correlation in demand can easily make a stochastic optimization approach intractable.

- *The solution is robust, protecting a risk-averse manager against adverse demand scenarios without being too conservative.*

The solution concept that motivates our model is the robust optimization approach. The robust approach protects the firm's revenue from very adverse outcomes, but is often criticized for being too conservative. With our sampling technique, we are able to use alternative regret-based robust objectives, which would be

intractable otherwise. Our numerical experiments show how the regret-based objectives can perform well in the lower tails of the revenue output distribution, without sacrificing much on the average performance. This advantage is particularly evident when the amount of data available is small.

- *We provide theoretical performance guarantees.*

Unlike the traditional robust pricing problem, the sampled problem is a convex program and can be efficiently solved. We prove that the sampled problem converges to the robust problem when the sample size increases. The design question that arises from this approach is how many samples do we need in order to have a performance guarantee on our sampling based solution. To answer this question, we define a notion of ϵ -robustness and show the sample size needed to achieve an ϵ -robust solution with some confidence level.

- *In the absence of enough real data, we generate data from an artificial distribution and still obtain theoretical guarantees.*

Data-driven approaches generally use the assumption that there is a large set of historical data available, which comes from the true underlying distribution of the uncertainty. This assumption can be quite restrictive in many real applications, for example when releasing a new product. For these cases, we introduce a new concept of random scenario sampling, where we use an artificial distribution over the uncertainty set to generate random sample points and apply the sampling based optimization framework.

The random scenario sampling framework we introduce is a rather powerful concept, given that we are now able to use a data-driven methodology to solve a robust problem without any actual historical data. Data-driven and robust optimization have been generally considered two separate fields, since they use very different initial assumptions of information about the problem. The bridge we develop between data-driven optimization and robust optimization has not been widely explored.

Additionally to the modeling contributions listed above, we also present a series of numerical experiments, where we study the simulated revenue performance of our dynamic pricing model. Using these numerical experiments, we demonstrate how the regret-based robust models can perform very well in practice, even compared to a sample-average benchmark. In particular, our models can be very useful when dealing with small sample sizes or when the sampling distribution is different from the true demand distribution.

Finally, we also show in this paper how to apply our methodology in practice, using a case study with actual airline data. Without considering many complexities of the airline industry, we show in a simplified setting how our pricing models might be used to optimize the revenues per flight.

1.2. Literature Review

A good introduction to the field of revenue management and dynamic pricing would include the overview papers of Elmaghraby and Keskinocak (2003), Bitran and Caldentey (2003) and the book by Talluri and

van Ryzin (2004). Most of the pricing literature uses the stochastic optimization framework, which relies on distributional assumptions on the demand model and often doesn't capture correlation of the demand uncertainty across time periods. Avoiding such problems, there are two different approaches prominent in the literature that will be most relevant for our research: data-driven and robust optimization. So far, these approaches have been studied separately and the modeling choice usually depends on the type of information provided to the firm about the demand uncertainty.

The operations management literature has explored sampling based optimization as a form of data-driven nonparametric approach to 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 Sample Average Approximation (SAA), when 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 and 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. Zhan and Shen (2005) apply the SAA framework for the single period price-setting newsvendor problem. Levi et al. (2007) also apply the SAA framework to the newsvendor problem (single and multi-period) and establish bounds on the number of samples required to guarantee with some probability that the expected cost of the sample-based policies approximates the expected optimal cost. Levi et al. (2012), using the assumption that the demand distribution is log-concave, develop a better bound on the number of samples to obtain a similar guarantee as in Levi et al. (2007). More specifically 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).

The fairly recent field of robust optimization proposes distribution-free modeling ideas for making decision models under uncertainty. This area was initiated by Soyster (1973) and it was further developed by Ben-Tal and Nemirovski (1998, 1999, 2000), Goldfarb and Iyengar (2003) and Bertsimas and Sim (2004). A robust policy can be defined in different ways. In this paper 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 in Scarf (1958), Gallego and Moon (1993), Ben-Tal et al. (2004), Bertsimas and Thiele (2006). The following papers by Adida and Perakis (2005); Nguyen and Perakis (2005); Perakis and Sood (2006); Thiele (2009); Birbil et al. (2006) are examples of the MaxMin robust approach applied to the dynamic pricing problem. This approach is usually appropriate for risk-averse managers, but it can give quite conservative solutions. For this reason we will explore the regret based models, which were originally proposed by Savage (1951). Lim and Shanthikumar (2007) and Lim et al. (2008) approach this

problem from a different angle, where the pricing policies are protected against a family of distributions bounded by a relative entropy measure. In the broader operations management literature, Yue et al. (2006) and Perakis and Roels (2008) use the MinMax Absolute Regret for the newsvendor problem. A comparison of MaxMin and MinMax Absolute Regret for revenue management can be found in Roels and Perakis (2010). An alternative approach is the relative regret measure, also known as the competitive ratio. In revenue management and pricing, Ball and Queyranne (2009) and Lan et al. (2008) use this MaxMin Ratio approach.

Irrespective of the demand uncertainty models described above, multi-period decision models can also be categorized between (i) closed-loop policies, where policies use feedback from the actual state of the system at each stage, and (ii) open-loop policies, where the entire policy is defined statically at the beginning of the time horizon. The initial robust framework discussed in the papers above does not allow for adaptability in the optimal policy. This open-loop robust framework has been applied to the dynamic pricing problem in Perakis and Sood (2006), Nguyen and Perakis (2005), Adida and Perakis (2005) and Thiele (2009). Moving towards 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, they advocate for the use of affinely adjustable policies. Recent work by Bertsimas et al. (2010) was able to show that optimality can actually be achieved by affine policies for a particular class of one-dimensional multistage robust problems. Unfortunately this is not our case, therefore we must admit that the affine policies we are using will only achieve an approximation to the fully closed-loop policy. Zhang (2006) develops a numerical study of the affinely adjustable robust model for the pricing problem using an MIP formulation. In this paper, we present a model that introduces an affinely adjustable approach to the dynamic pricing problem and uses sampling based approach to solve the robust problem.

As mentioned by 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 that arises is how many scenarios do we need to have any confidence guarantees on our model's solution. To answer this question, Calafiore and Campi (2005, 2006) define the concept of an ϵ -level robust solution and provide a theoretical bound on the sample size necessary to obtain this solution. The bound was later improved in Campi and Garatti (2008) and Calafiore (2009), which is provably the tightest possible bound for a class of robust problems defined as "fully-supported" problems. Recently Pagnoncelli et al. (2009) suggested how to use this framework to solve chance constrained problems.

1.3. Outline

The remainder of the paper is structured as follows. In Section 2, we introduce our modeling approach and develop the theoretical performance guarantees. In Section 3, we show the simulated performance of the proposed models and interpret the numerical results. In Section 4, we develop a case study with actual

airline data. Finally, in Section 5, we conclude with a summary of the theoretical and numerical results. The appendices provide supplementary content to the reader that were omitted from the paper for conciseness. Appendix A provides a summary of the notation used. Appendices B-E display the proofs of our theoretical results in Section 2.

2. The Model

Before introducing the general model we propose in this paper, we motivate the problem with the following example. Suppose a firm sells only one product over a two period horizon, with a limited inventory of C . Moreover, suppose the firm has a set of N historical data samples of demand and prices for each period of the sales horizon. We will assume for this example the demand is a linear function of the price plus δ , which is a random noise component: $\text{Demand}_t = a_t - b_t \text{Price}_t + \delta_t$. After estimating the demand function parameters (a_1, a_2, b_1, b_2) using the N data points, we are left with a set of estimation errors (i.e. the difference between the realized demand data and the estimated demand): $\delta_t^{(1)}, \dots, \delta_t^{(N)}$ for each time period $t = 1, 2$. 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. It is not clear how one should define this uncertainty set given a pool of uncertainty samples and the resulting problem can also 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 paper. Our goal, as seen in the following model, is to find a pricing strategy that robustly maximizes the firm's revenue with respect to the N given observations of demand uncertainty:

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

Note that the given uncertainty samples $\delta_t^{(i)}$ will approximate the uncertainty set from the traditional robust optimization approach. The major theoretical challenge now is to determine how many samples are needed for the sampling based model to approximate the original robust problem.

On the modeling aspect of the problem, one of the main problems with the solution concept presented above is that this MaxMin robust approach can often be too conservative. For this reason, we will propose other types of robust modeling. Another problem in this example is that the second period price p_2 does not depend on the uncertainty realized on the first period δ_1 , which is what we call an open-loop model. Ideally, the second period pricing policy should be a function of the new information obtained in the first period, $p_2(\delta_1)$, which is known as a closed-loop model.

After the motivating example illustrated above, we proceed to generalize the problem to include components such as salvage/overbooking inventory, nonlinear demand models, and regret-based objective functions.

Throughout Section 2.1 we develop a solution approach that mitigates the modeling issues described before. Further on, in Section 2.2, we will address the theoretical issue of the number of samples required. We refer the reader to Appendix A for a summary of the notation that will be used in this paper.

2.1. Model definition

Let T be the length of the time horizon. Define p_t as the price at time t and the nominal demand function as d_t . The nominal demand d_t captures the deterministic part of the demand for a given price level. It is important to note that the modeling techniques and the results presented in this paper 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.

It is also important to note that the model presented in this paper can be easily extended to a multi-product pricing problem, with multiple resources. The results, both theoretical and numerical, would remain the same. Therefore, this extension is not explored in this paper to avoid complicating the notation without adding any additional insight.

For technical reasons, which are discussed later, we require that the demand function satisfies the following convexity/concavity assumption.

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

The intuition behind Assumption 1 is that we want the space of pricing strategies to be a closed convex set and the objective function to be strictly concave, giving rise to a unique optimal solution. Examples of demand functions that satisfy Assumption 1 which are common in the revenue management literature are the linear, iso-elastic, and the logarithmic demand functions.

To capture the variability in demand, given the nominal demand function, $d_t(p_t)$, 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, we would have: $\tilde{d}_t(p_t, \delta_t) = d_t(p_t) + \delta_t$. In the literature, it is mostly 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 prices p_t are defined as a function of control variables s and the uncertainty δ , therefore denoted $p_t(s, \delta)$. We also refer to s as the pricing policy. We assume there is a finite set of decision variables s , which lie within the strategy space S . We assume that S is a finite dimensional and compact set. In the case of static pricing (open-loop policies), s is a vector of fixed prices decided before hand, independent of the realizations of demand. 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 will express the actual realized prices as $p_t(s, \delta)$ for both cases. For an example of this pricing policy, see Section 2.2, equation (6). Also to avoid complicating the notation, 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 determine the price at time t , therefore also determining the realized demand $\tilde{d}_t(s, \delta)$. In this paper, we restrict ourselves to the family of pricing functions $p_t(s, \delta)$ that is affine in s . This restriction is formalized later in Assumption 2.

In general, $\delta = (\delta_1, \dots, \delta_T)$ is a random vector with one 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 about 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 nonnegative and the total demand seen by the firm should be less than its total capacity C or else the firm will pay an overbooking fee o for every unit sold above capacity. For every unit of capacity not sold, the firm will get 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. Define $w(s, \delta)$ as the terminal value of remaining inventory, which can be either an overbooking cost or a salvage value revenue. The terminal inventory is the difference between the capacity and the number of units sold: $C - \sum_{t=1}^T \tilde{d}_t(s, \delta)$. Define the terminal value as:

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

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

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

As stated before, our goal is to find a pricing policy that will give a robust performance for all possible realizations of demand. One can think of the robust pricing problem defined 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 to find the best robust policy under the constraints that the pricing policy yields nonnegative prices and that the total demand must be less than or equal the capacity (although relaxed by overbooking fees and salvage value). To express the different types of robust objectives explored in this paper, define $h^{obj}(s, \delta)$ as the objective function realization for a given pricing strategy s and uncertainty δ . The index *obj* can be replaced with one of three types of robust objectives that we consider in this work: the **MaxMin**, the **MinMax Regret** and the **MaxMin Ratio**. For example in the MaxMin case, the objective function is simply given by $h^{MaxMin}(s, \delta) = \Pi(s, \delta)$. The remaining two objectives will be further explained later.

The following model defines the robust pricing model.

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

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

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

The model above for the hindsight revenue $\Pi^*(\delta)$ is a deterministic convex optimization problem, which can be efficiently computed for any given δ .

Define the absolute regret as the difference between the hindsight revenue and the actual revenue: $\Pi^*(\delta) - \Pi(s, \delta)$. To be consistent with the MaxMin formulation, we can define the objective function for the MinMax Regret as $h^{Regret}(s, \delta) = \Pi(s, \delta) - \Pi^*(\delta)$, which is the negative of the regret. We will continue calling this MinMax Regret, since it is a more common term in the literature (although this would be more precisely named MaxMin Negative Regret). Finally, the MaxMin Ratio, which is also known as the relative regret or competitive ratio, tries to bound the ratio between the actual revenue and the hindsight revenue. The objective function for the MaxMin Ratio can be concisely written $h^{Ratio}(s, \delta) = \frac{\Pi(s, \delta)}{\Pi^*(\delta)}$.

The inputs to the general robust model are: the structure of the pricing policy $p_t(s, \delta)$, parameterized by the decision variables s ; the demand functions $\tilde{d}_t(s, \delta)$ for any given price policy s and uncertainty δ ; the objective function $h^{obj}(s, \delta)$ and the uncertainty set U (which we will later replace with uncertainty samples from U , as we can see in Section 2.2). The outputs are the set of pricing decisions s and the variable z , which is a dummy variable used to capture the robust objective value of the pricing policy.

Note that when using adjustable policies, the profit function $\Pi(s, \delta)$, which defines the objective function, is neither concave nor convex with respect to the deviations δ . In Appendix B, we illustrate this with a simple instance of the adjustable MaxMin model. Because of this lack of convexity or concavity, the traditional robust optimization methods (i.e. solve the exact robust problem using duality arguments or simply searching over the boundary of uncertainty set) will be intractable. Note the example in Appendix B uses a simple linear demand model and simple objective function. More complex models will generally still have this non-convexity issue. In the next section we introduce the sampling based approach that we advocate for solving the robust pricing problem.

2.2. Sampling based optimization

Ideally, we would like to solve the exact robust pricing problem, but as we have seen in the previous section this can easily become 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 (3) by a finite number of constraints. It is only natural to question how good is the solution to the sampling-based problem relative to the original robust problem that was proposed in the previous section. In order to discuss this issue, we will explore in this section the theoretical framework needed to analyze the robustness of a sampling-based solution. For clarity purposes, we again refer the reader to Appendix A for a summary of the notation used in this section to formulate the theoretical results. Define the following model as the sampling based counterpart of the robust pricing model:

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

The idea of using sampled uncertainty scenarios, or data points, in stochastic optimization models is often called sampling or scenario based optimization. For tractability, we will make the following linearity assumption about the set of pricing policies that we admit.

Assumption 2. *We will restrict our pricing function $p_t(s, \delta_1, \dots, \delta_{t-1})$ to policies that are linear 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 the dependence on δ can be nonlinear. To illustrate this affine family of functions, the following is the pricing function used in our numerical studies:

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

In words, $s_{1,t}$ is the static component of price and $s_{2,t}$ adjusts the price linearly according to the cumulative deviation from the expected sales level 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 from Assumptions 1 and 2, the constraints in (5) define a convex feasible set for any fixed vector δ . Therefore the scenario based problem is a convex optimization problem and can be solved by any nonlinear optimization solver. It is easy to argue that the exact robust problem that we initially stated in (3) has an optimal solution. For a proof of Proposition 1, see Appendix Appendix C.

Proposition 1. *There exist an optimal solution to (3).*

It remains to show how good is the approximation of the scenario based model relative to the exact robust model. Depending on what type of demand information and historical data that is initially provided to the firm, we propose two solution approaches: the Data-Driven and the Random Scenario Sampling. The first approach, Data-Driven, assumes 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 don't 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.

Suppose we are given a large sample set of historical data of prices and demands. There are many ways to estimate the parameters of the demand function using such data. As an example, in the numerical study of Section 4, we used linear regression on the price-demand data to estimate a linear demand model. The estimation error obtained is the sample deviations δ in our model. Under the assumption that the data points come from the true underlying distribution, we will show a performance guarantee on how robust is the sampling based solution relative to the exact solution.

In order to simplify notation, define the decision variable $x = (s, z)$ as the combined pricing policy decision s and objective function value z . Note that it lies within the domain X , i.e. $x \in X$, 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 define the pricing problem using a standard convex optimization form with a linear objective. Define the equivalent constraint function with a scalar valued g such that: $g(x, \delta) = \max \left\{ z - h^{obj}(s, \delta), -p(s, \delta) \right\}$. Since each constraint in the definition of $g(x, \delta)$ above is convex in x for any given δ , the maximum between them is still convex in x . Moreover, $g(x, \delta) \leq 0$ is equivalent to the constraint set defined before in (3). Then problem (3) can be concisely defined as the following model (7):

$$\max_{x \in X} c'x, \quad \text{s.t.} \quad g(x, \delta) \leq 0, \quad \forall \delta \in U \quad (7)$$

Since we cannot solve the problem with a continuum of constraints, we solve the problem for a finite sample of deviations $\delta^{(i)}$ from the uncertainty set U , where $i = 1, \dots, N$. Then the concise version of the sampled problem (5) can be defined as the following model (8):

$$\max_{x \in X} c'x, \quad \text{s.t.} \quad g(x, \delta^{(i)}) \leq 0, \quad \forall i = 1 \dots N \quad (8)$$

To validate the data-driven approach presented in (8), we provide a convergence analysis result. We show that the optimal solution of the sampled problem (8) converges to the robust solution of (7) when the number of samples grows to infinity. The analysis then justifies from a theoretical point of view the validity of a sampling based approximation solution. The following assumptions are needed for the convergence proof. Denote the feasible region of the robust problem (7) as \bar{X} , i.e., $\bar{X} = \{x \in X : g(x, \delta) \leq 0, \delta \in U\}$. Similarly, let $\bar{X}^{(N)} = \{x \in X : g(x, \delta^{(i)}) \leq 0, i = 1, \dots, N\}$ denote the feasible region of (8).

Assumption 3. (i) *The feasible region \bar{X} of (7) 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 is denoted by f , which is bounded below: there exists $\varepsilon > 0$ such that $f(\delta) > \varepsilon$ for all $\delta \in U$.

Define the optimal value of (7) by Z^* and the optimal value of (8) by $Z^{(N)}$. 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 optimal solution $x^{(N)}$ and the optimal value $Z^{(N)}$ converge almost surely to the optimal solution x^* and the optimal value Z^* , respectively.*

Proof. We prove the convergence result of Theorem 1 by the three following steps.

Step 1: The sequence $\{\delta^{(i)}\}$ is almost surely a dense subset of U .

Let $\delta \in U$, and let $\gamma > 0$. We want to show that there exists i such that $\|\delta^{(i)} - \delta\| < \gamma$. Let $(\hat{\delta}^{(j)})$ be a dense sequence in U . Since $(\hat{\delta}^{(j)})$ is a dense sequence in U , there exists j such that $\|\delta - \hat{\delta}^{(j)}\| < \gamma/2$. Next, by Assumption 3 (iii), with probability 1, there exists i such that $\|\delta^{(i)} - \hat{\delta}^{(j)}\| < \gamma/2$.

Now, we define for each j the event $A_j = \{\exists i \text{ s.t. } \|\delta^{(i)} - \hat{\delta}^{(j)}\| < \gamma/2\}$ and we have $\mathbb{P}(A_j) = 1$ for all j . Therefore, by using Assumption 3 (ii) and the Borel-Cantelli Lemma, we obtain: $\mathbb{P}(\cap_j A_j) = 1$. In other words, $\mathbb{P}(\forall j, \exists i \text{ s.t. } \|\delta^{(i)} - \hat{\delta}^{(j)}\| < \gamma/2) = 1$.

Finally, by using the triangle inequality, we obtain: $\mathbb{P}(\forall \delta \in U, \exists i \text{ s.t. } \|\delta^{(i)} - \delta\| < \gamma) = 1$. Taking all rational γ and intersecting the events $\{\forall \delta \in U, \exists i \text{ s.t. } \|\delta^{(i)} - \delta\| < \gamma\}$, we conclude that the sequence $(\delta^{(i)})$ is almost surely a dense subset of U .

Step 2: The feasible regions $\bar{X}^{(N)}$ form a decreasing sequence and $\bar{X}^{(N)} \downarrow \bar{X}$ almost surely.

$\bar{X}^{(N)}$ is a decreasing sequence since we are adding more constraints and then restricting the feasible region. Let $Y = \lim_{N \rightarrow \infty} \cap_N \bar{X}^{(N)}$ and suppose by contradiction that $Y \neq \bar{X}$. Note that $\bar{X} \subset \bar{X}^{(N)}$ for all N and hence $\bar{X} \subset Y$.

Let $x \in Y \setminus \bar{X}$. Since $x \notin \bar{X}$, there exists $\delta \in U$ such that $g(x, \delta) > 0$. By the continuity of the function g (see proof of Proposition 1), there exists an open neighborhood W around δ such that for all $\delta' \in W$, $g(x, \delta') > 0$. We have shown that the sequence $\delta^{(i)}$ is almost surely a dense subset of U and then there exists i such that $\delta^{(i)} \in W$ and $g(x, \delta^{(i)}) \leq 0$ (almost surely). Therefore, this is a contradiction and $Y = \bar{X}$ so that $\bar{X}^{(N)} \downarrow \bar{X}$ almost surely.

Step 3: $\bar{X}^{(N)}$ is eventually a compact set.

Let $C^{(N)}$ denote the recession cone of $\bar{X}^{(N)}$, i.e., $C^{(N)} = \{y \in \mathbb{R}^{\dim(x)} : \forall x \in \bar{X}^{(N)}, \forall \lambda \geq 0 : x + \lambda y \in \bar{X}^{(N)}\}$. We know that $C^{(N)}$ is a decreasing sequence and that: $\cap_N C^{(N)} = C_{\cap_N \bar{X}^{(N)}} = C_{\bar{X}} = \{0\}$ (the first equality follows because $\bar{X}^{(N)}$ are closed convex sets, whereas the second equality holds due to Assumption 3 (i)). Now let S be the unit sphere and consider $C^{(N)} \cap S$. This is a decreasing sequence as well and $\cap_N (C^{(N)} \cap S) = (\cap_N C^{(N)}) \cap S = \{0\}$. Since $\bar{X}^{(N)}$ are closed by definition, so are $C^{(N)}$ and hence $C^{(N)} \cap S$ are compact.

$C^{(N)} \cap S$ is a decreasing sequence of compact sets. Therefore, if all of them are non-empty their intersection has to be non-empty. We then conclude that at least one of the $C^{(N)} \cap S$ is empty and then all the subsequent ones are empty as well. This shows that $\bar{X}^{(N)}$ is eventually a compact set.

We are now in position to complete the proof of Theorem 1. Let the optimal solution of (8) be $x^{(N)}$. Since $\bar{X}^{(N)}$ is eventually a compact set, the sequence $x^{(N)}$ has a limit point, say x^* . Since $\bar{X}^{(N)}$ are decreasing, the limit point $x^* \in \bar{X}^{(N)}$ for all N and hence $x^* \in \bar{X}$. Therefore, $c'x^*$ is a limit point of $c'x^{(N)} = Z^{(N)}$, and since $Z^{(N)}$ is a decreasing sequence, $Z^{(N)} \downarrow c'x^*$. However, x^* is a feasible solution for (7) and $c'x^* \leq Z^* \leq Z^{(N)}$ for all N . By taking the limit when $N \rightarrow \infty$ in both sides, we obtain that $Z^* = c'x^*$. We then conclude that $Z^{(N)} \rightarrow Z^*$ and $x^{(N)} \rightarrow x^*$ almost surely. \square

We note that the result of Theorem 1 implies that a series of convex problems converge to a non-convex problem with probability 1. This result appears to be interesting and validates the approach of using the sampled problem to approximate its robust counterpart. At this point, we showed that the sampled problem will converge to the robust problem with infinite samples. For any practical application, it is also important to understand what will happen to the sampled solution with a finite number of samples. For that purpose, we obtain a bound on how close your sampled solution will be to the robust solution. The following definition of ϵ -robustness is required to develop this bound.

Definition 1. For a given pricing policy x and a distribution \mathcal{Q} of the uncertainty δ , 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 case, given the pricing policy $x = (s, z)$, $V_{\mathcal{Q}}(x)$ is the probability that the actual realization of demand gives the firm a revenue lower than z , which we computed as the worst-case revenue, or that it violates non-negativity constraints. In reality, the constraints are “naturally” enforced. The firm won’t set a negative price, so if such a deviation occurs, the constraints will be enforced at a cost to the firm’s revenue. Therefore, it is easy to understand any violation as an unexpected loss in revenue. We can now define the concept of ϵ -robust feasibility.

Definition 2. We say x is ϵ -level robustly feasible (or simply ϵ -robust) if $V_{\mathcal{Q}}(x) \leq \epsilon$.

Note that the given set of scenario samples is itself a random object and it comes from the probability space of all possible sampling outcomes of size N . For a given level $\epsilon \in (0, 1)$, a “good” sample is one such that the solution $x^N = (s^N, z^N)$ to the sampling based optimization model will give us an ϵ -robust solution, *i.e.*, 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 used N .

Before we introduce the main result, there is one last concept that needs to be explained. Suppose that we do not have samples obtained from the true distribution \mathcal{Q} , i.e. we do not have enough historical data. Instead we are given the nominal demand parameters and the uncertainty set U . We would like to be able to draw samples from another chosen distribution \mathcal{P} and run the sampling based pricing model (8). In order to make a statement about the confidence level of the solution and the sample size, we must make 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 the distribution \mathcal{P} from which we draw the samples. 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 (which is usually the case in data-driven problems). Note that the assumption above will be satisfied under a more restrictive, but perhaps more common, assumption for continuous distributions of Bounded Likelihood Ratio $\frac{d\mathbf{P}_{\mathcal{Q}}(x)}{d\mathbf{P}_{\mathcal{P}}(x)} \leq k$.

At first glance, it seems hard for a manager to pick a bound k on the likelihood ratio that would work for his uncertainty set and sampling distribution without any knowledge of the true underlying distribution. On the other hand, the variance of the demand distribution is a familiar statistic to most managers and might be somehow obtained by the firm. Gaur et al. (2007) propose a way to estimate the variance of the demand distribution using the dispersion of managers' judgemental forecasts. With the variance of the demand distribution and the additional assumption that the demand uncertainty is independent across time periods, Lobel (2012) proposes a way to derive a likelihood ratio bound k between a uniform distribution and any log-concave distribution. Similar results can also be obtained by using other statistics about the volatility of the demand. Also note that the family of log-concave distributions, as defined by distributions where the log of the density function is concave, is a rather extensive family. For the remainder of the paper, we will assume that such bound k is known by the manager.

Let n_x be the dimension of the strategy space X , x^N be the solution of (8) using the N sample points, and ϵ be the robust level parameter. The following theorem develops a bound on the probability that the solution of the sampled problem is not ϵ -robust, i.e. probability of drawing a "bad" sample.

Theorem 2. Assume the sampling distribution \mathcal{P} bounds the true uncertainty 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 $(1 - \beta(N, \epsilon))$, the solution x^N is ϵ -level robustly feasible, i.e.,

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

In other words, the level $\beta(N, \epsilon)$ is a bound on the probability of getting a “bad” sample of size N for a robust level ϵ . Then $1 - \beta(N, \epsilon)$ is the confidence level that our solution is ϵ -robust. For proof of Theorem 2, see Appendix Appendix D.

As a corollary of Theorem 2, we can obtain a direct sample size bound for a desired confidence level β and robust level ϵ .

Corollary 1. *If the sample size N follows:*

$$N \geq N(\epsilon, \beta) \doteq (2k/\epsilon) \ln(1/\beta) + 2n_x + (2n_x k/\epsilon) \ln(2k/\epsilon)$$

Then with probability greater than $1 - \beta$ the solution to the sampled problem will be ϵ -level robustly feasible.

The bound in Corollary 1 is not necessarily the tightest value of N that will satisfy Theorem 2 for a given β and ϵ . Numerically solving for N the equation $\beta = \binom{N}{n_x} (1 - \epsilon/k)^{N - n_x}$ might give a smaller sample size requirement. On the other hand, Corollary 1 offers a direct calculation and insight about the relationship between the sample size and the confidence/robust level of the sampled solution. Note that $N(\epsilon, \beta)$ goes to infinity as either ϵ or β go to zero, which is rather intuitive. On the other hand the dependence on β is of the form $\ln(1/\beta)$ which means that the confidence parameter β can be pushed down towards zero without significant impact on the number of samples required. For implementation purposes, it allows us to keep a good level of confidence β and design N based on the ϵ -level of robustness desired.

A key difference between the proof of Theorem 2 and Theorem 1 in Calafiore and Campi (2006) is the introduction of randomized sampling. The latter could not be directly applied in our case since we generalized the input data to allow for the random sampling approach. Instead, we had to introduce the bounded likelihood ratio assumption to connect the true distribution to the sampling distribution.

In order to apply this confidence bound to our pricing problem, we only need to know the likelihood bound k and the dimension of the strategy space n_x , in other words, the number of decision variables in the optimization model. The latter will depend on the functional form of the pricing policy chosen. 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: $n_x = T + 1$. If the pricing policy is adjustable, it requires more decision variables. In particular, for a price that is a linear function of the sum of previously realized uncertainties, which we use in Section 3, $n_x = 2T$. For instance, in the airline case study of Section 4, we deal with a 3-period adjustable model, therefore $n_x = 6$.

Note that this Random Scenario Sampling result can be quite useful in practice, specially if there is a small amount of data or no data samples 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.

3. Numerical Results

In this section we present simulation results using the models described before. All optimization problems discussed below were solved using AMPL modeling language and LOQO solver.

The following experiments are designed to compare the out-of-sample performance of the multiple pricing objectives discussed in Section 2 in order to obtain insights about the resulting pricing policies. More specifically, we will compare the following models: MaxMin, MinMax Regret, MaxMin Ratio, and Sample Average Approximation. Note that we are introducing here the Sample Average Approximation (SAA) approach only as a benchmark to our robust models and we will explain it 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 3.2.

In order to isolate the differences between the objective functions above, we will reduce the problem to a linear demand function with additive uncertainty. In particular, the set of parameters used in this section are: $T = 2$, $d_t(p_t) = 200 - 0.5p_t$, $C = 120$, $o = 1000$, $v = 10$. These numbers are chosen to be in a similar order of magnitude to the case study in Section 4. For simplicity, we will also consider a box-type uncertainty set, where $U = \{\delta : |\delta_t| \leq 15, \forall t = 1, 2\}$. The reader should note that this demand model is used here for simplicity of notation only and the use of other demand functions and uncertainty models can be easily incorporated. In Sections 3.1 and 3.2, we assume the uncertainty to be uniformly distributed over the entire uncertainty set and we sample the data points directly from this distribution (Data Driven). In Section 3.3, we still use the uniform distribution for sampling and optimization purposes, but define the true underlying distribution to be a truncated normal distribution that will be explained further later (Random Scenario Sampling). The adjustable pricing policy used is displayed in Section 2.2, equation (6). We have performed a variety of alternative experiments, with different demand functions, sampling distributions, pricing functions and a multi-product extension. The qualitative insights with these alternative experiments were the same, therefore we will not display them in this paper.

3.1. Testing the model: illustrating the output revenue distribution

Define the **MaxMin** model as in the traditional robust optimization model, which maximizes the worst case revenue $\Pi(s, \delta^{(i)})$ for each data point $\delta^{(i)}$. For the regret-based models, we can solve the hindsight revenue $\Pi^*(\delta^{(i)})$ for each data point $\delta^{(i)}$ by solving a simple convex optimization problem. The **MinMax Regret** is defined by subtracting the hindsight term from the actual revenue: $\Pi(s, \delta^{(i)}) - \Pi^*(\delta^{(i)})$. Similarly, the **MaxMin Ratio** is defined by dividing the actual revenue by the hindsight revenue: $\Pi(s, \delta^{(i)}) / \Pi^*(\delta^{(i)})$.

Robust optimization is often criticized for generating very conservative policies. In other words, it provides good performance in the worst-case revenue (or regret/competitive ratio), but bad overall average performance. For this reason, we will compare the solution of our robust models with a more traditional stochastic optimization framework, where we want to maximize the expected value of the revenue. Again, we use the pricing decision as a function of the previous realized uncertainty. In this case, the objective approximates the expected value of the revenue by taking the average of each revenue point: $h^{SAA}(s, \delta^i) = \frac{1}{N} \sum_{i=1}^N \left(\sum_{t=1}^T p_t(s, \delta^i) \tilde{d}_t(s, \delta^i) + w(s, \delta^i) \right)$. This solution concept is what we refer to as Sample Average Approximation (**SAA**).

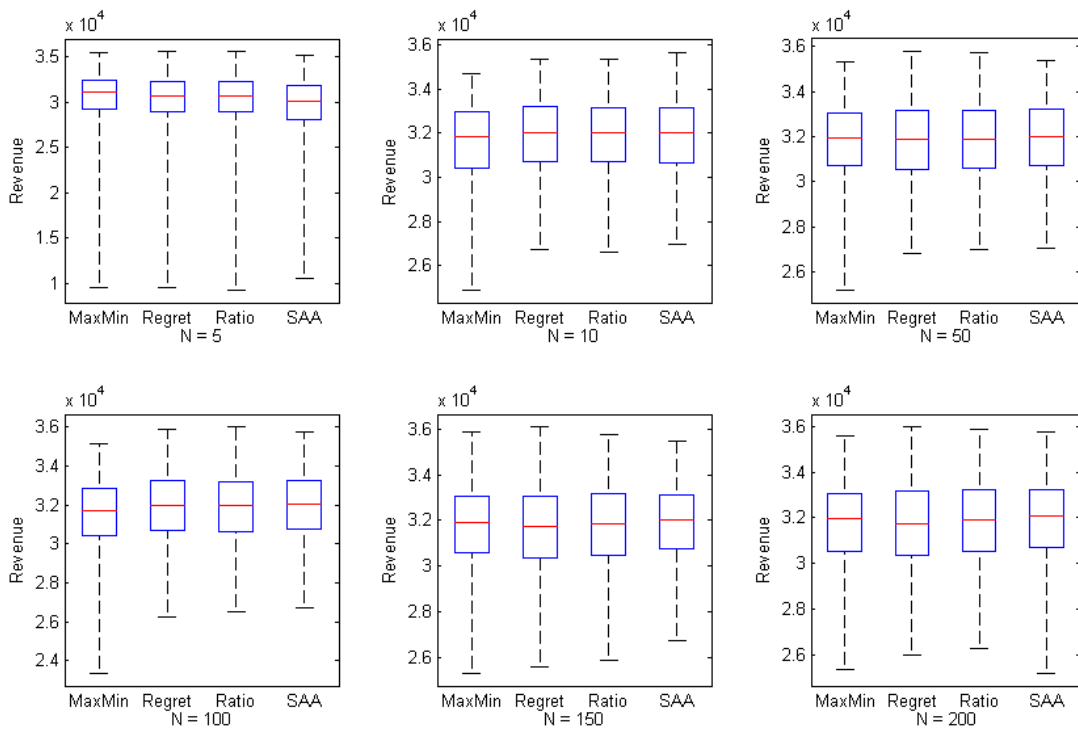


Figure 1: Example of out-of-sample revenue distribution for single product case. Sample sizes vary from 5 to 200 samples.

In Figure 1, we observe the out-of-sample revenue performance of our four proposed models with different sample sizes varying from 5 to 200. The revenue ‘boxplots’ displayed throughout this section can be explained as follows: The middle line indicates the median; The boxes contain 50% of the observed points (25% quantile to the 75% quantile); The lines extending from the boxes display the range of the remaining upper 25% and lower 25% of the revenue distribution. *Not much can be concluded from this particular instances’ revenue snapshot. It illustrates how the intrinsic volatility of the output revenues creates too much noise for us to make precise qualitative judgements between the different models.* This issue can be fixed by performing

paired tests between different models, which compares the difference in revenues between models taken at each data sample. In these paired tests displayed in the following sections, the intrinsic volatility of revenues is reduced by the difference operator, which allows us to compare the models' performances.

3.2. Pairwise comparing the models

For the remainder of this section we will perform pairwise comparisons between the MaxMin Ratio and the other models, namely the SAA, MaxMin and MinMax Regret. The reason we focus on the MaxMin Ratio model is because it demonstrated to be the most balanced between the robust models and the SAA, as shown in the next experiments. The other pairwise combinations were also experimented, but are not displayed in this paper as they don't provide any further insight.

To conduct the pairwise comparisons, we measure the average difference in revenues and construct an empirical confidence interval in order to draw conclusions about which model performs better on the average revenue. Another useful statistic to measure robust models is known as Conditional Value-at-Risk (CVaR), which is also known in finance as mean shortfall. Taken at a given percentile α , the 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. 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, define for the positive side of the distribution, $\alpha > 50\%$, the CVaR_α to be the expected value of a variable conditional on it being above the 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 , 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 issue to be observed is how the different models behave when provided small sample sizes versus large sample size. For this reason, we will perform the experiments in the following order, starting with $N = 5$ and increasing until $N = 200$:

1. Take a sample of size N uncertainty data points;
2. Optimize for different models to obtain pricing policies;
3. Measure out-of-sample revenues for 2000 new samples, keeping track of pairwise differences in revenues;
4. Record statistics of average difference and CVaR of the difference in revenues for each pair of models;
5. Repeat from steps 1-4 for 1000 iterations to build confidence intervals for the average difference and CVaR of the difference;
6. Increase the sample size N and repeat 1-5.

3.2.1. Comparing SAA and MaxMin Ratio

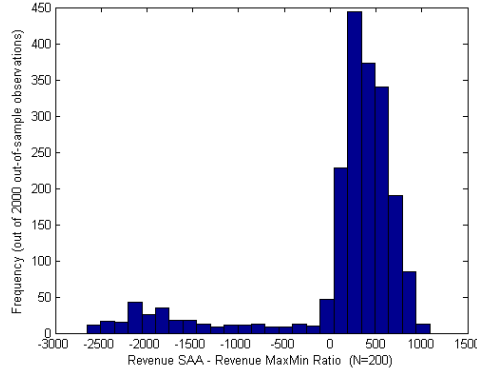


Figure 2: Comparison between SAA and MaxMin Ratio

In the first test we compare the revenues of the SAA and the MaxMin Ratio model. In Figure 2, we display a histogram of the difference of out-of-sample revenues for a particular instance of the single product problem, using $N = 200$ samples. More specifically, for each of the 2000 out-of-sample revenue outcomes generated, we calculate the difference between the revenue under SAA policy minus the revenue under MaxMin Ratio model. By calculating the statistics of this histogram, we obtain a mean revenue difference of 133 and standard deviation of 766. In fact, we observe there is an indication that the SAA obtains a small revenue advantage on a lot of cases, but can perform much worse in a few of the bad cases.

We also obtained $CVaR_{5\%} = -2233$ and the $CVaR_{95\%} = 874$. One can clearly 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 quite limited.

Up to this point, 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 for these estimates. We are also interested in finding out if these differences appear 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 $CVaR_{5\%}$ of this difference is significantly larger than $CVaR_{95\%}$ (in absolute terms).

In Table 1, we first display the mean average difference in the SAA revenue minus the MaxMin Ratio revenue, with a 90% confidence interval (third column) calculated from the percentile of the empirical distribution formed over 1000 iterations. One can initially observe that for $N = 150$ and $N = 200$ the average difference is clearly positive. In other words, *with large enough sample size, the SAA performs better on average than the MaxMin Ratio.*

When N is reduced, the average differences move closer to zero. When $N = 5$ and $N = 10$, the average difference in revenues between the SAA and the MaxMin Ratio is in fact negative, although not statistically significant since the confidence interval contains zero. This is implying that *when the sample size is small the MaxMin Ratio will often perform better on average.*

N	Average		5% CVaR		95% CVaR	
	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-79	[-1763,909]	-2686	[-11849,-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]

Table 1: Average and CVaR of the difference: Revenue SAA - Revenue MaxMin Ratio

In Table 1, we also present the lower 5% and upper 95% tails of the revenue difference distribution. We can observe in this case that when the sample size is small, the two sides of the distribution behave rather similarly. As N increases, we can observe the lower tail of the revenue difference distribution is not very much affected and the length of the confidence interval seems to decrease around similar mean values of $CVaR_{5\%}$. The upper tail of the distribution, on the other hand, seems to shrink as we increase N . When $N = 200$, the upside benefits of using SAA rather than MaxMin Ratio will be limited by 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 $CVaR_{5\%}$ will be larger in absolute value than the $CVaR_{95\%}$, since the 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 paired comparison between the SAA and the MaxMin Ratio, we observed 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 testing for these results can't either confirm or deny this conclusion. For large sample sizes, the results are more clear: with 90% confidence level, the SAA will obtain a better average revenue than the MaxMin Ratio, but will pay a heavy penalty in the worst revenue cases. This last conclusion can be viewed as intuitive, since the SAA model tries to maximize average revenues, while the robust models like the MaxMin Ratio try to protect against bad revenue scenarios. On the other hand, the better performance of the robust model over the SAA using small sample sizes is not intuitive and understanding this behavior leads to interesting directions for future research.

Before proceeding to the next experiments among the robust models, we should make a note that

similar paired tests of the the SAA with the other robust models (MaxMin and MinMax Regret) were also performed and produced similar results as the ones displayed above using the MaxMin Ratio. In the interest of conciseness, we decided not to display them here as they would not add any additional insight.

3.2.2. Comparing MaxMin Ratio and MaxMin

In the following experiments we will perform paired comparisons between the different robust models. In Figure 3, we display the histogram of the paired comparison between the revenues of the MaxMin Ratio and the MaxMin models. This histogram displays the distribution over 2000 out-of-sample data points in the difference in revenues for one particular iteration of the experiment. Although the distribution seems skewed towards the negative side, the average is actually positive, due to the long tail in the right-hand side. We would like to know for multiple iterations of the experiment if this average remains positive and if the tail of the distribution is much larger on the positive side, as displayed in this histogram. Moreover, we would like to compare these results for multiple initial sample sizes.

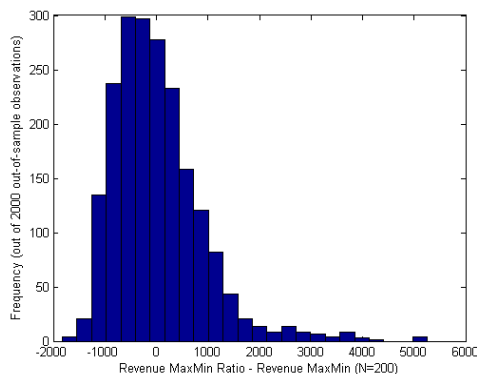


Figure 3: Comparison between the MaxMin Ratio and the MaxMin

The results of Table 2 show that the MaxMin Ratio does have a consistently better average performance than the MaxMin model for all sample sizes tested, but it is a small margin and the confidence intervals cannot support or reject this claim. Also, we can see that the MaxMin Ratio is more robust than the MaxMin, as the 5% best cases for the MaxMin Ratio (95% CVaR column) are significantly better than the 5% best cases of the MaxMin (5% CVaR column). Note that for $N \geq 50$ the confidence intervals for the 95% and 5% CVaR do not intersect, suggesting that with a 90% confidence level, the MaxMin Ratio is in fact more robust than the MaxMin under this 5% CVaR measure.

In summary, *the MaxMin Ratio model outperforms the MaxMin in all areas*. It displayed better average revenues and shortfall revenues (5% CVaR) than the MaxMin for all sample sizes. It is rather counter intuitive that the MaxMin Ratio model can be at the same time more robust and have better average

N	Average		5% CVaR		95% CVaR	
	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	577	[-322,2454]	-1369	[-4072,-245]	3912	[997,10884]
10	341	[-202,1126]	-1178	[-2847,-400]	3382	[1556,7367]
50	80	[-266,472]	-1149	[-1924,-614]	3267	[2014,5132]
100	50	[-200,381]	-1125	[-1619,-655]	3025	[1914,4116]
150	21	[-206,310]	-1109	[-1522,-687]	3031	[1981,4079]
200	23	[-204,289]	-1107	[-1481,-713]	2934	[1848,3925]

Table 2: Average and CVaR of the difference: Revenue MaxMin Ratio - Revenue MaxMin

performance. This could be a result of the sampling nature of these models. More specifically, the worst case revenues for the pricing problem are usually located in the borders of the uncertainty set, which is a hard place to estimate from sampling. When using regret based models, like the MaxMin Ratio, we adjust the revenues by the hindsight revenues. This shift in the position of the worst-case deviations of demand can make it easier for a sampling procedure to work.

3.2.3. Comparing MaxMin Ratio and MinMax Regret

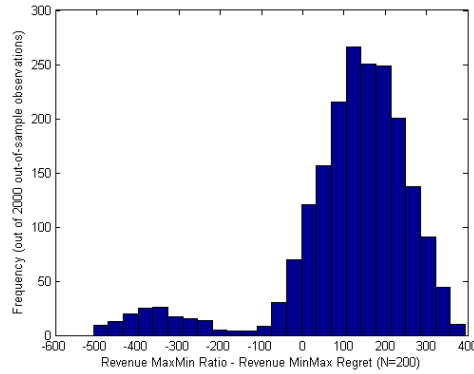


Figure 4: Comparison between MaxMin Ratio and MinMax Regret

In the next experiment, we compare the two regret-based robust models: MaxMin Ratio and MinMax Regret. Figure 4 displays the difference distribution in revenues from the MaxMin Ratio and MinMax Regret for one particular iteration of the experiment. We can see that there is a strong bias towards the positive side, showing that the MaxMin Ratio performs better on average than the MinMax Regret, but there is a small tail on the left with much higher values. When observing the difference in revenues of these two

models for 1000 iterations, we obtain the following Table 3 of statistics.

N	Average		5% CVaR		95% CVaR	
	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-16	[-97,32]	-232	[-796,-25]	106	[10,321]
10	-13	[-90,41]	-341	[-2847,-91]	138	[29,350]
50	15	[-48,100]	-483	[-1924,-231]	296	[76,415]
100	54	[-6,140]	-458	[-1619,-267]	240	[94,386]
150	68	[9,159]	-456	[-1522,-292]	260	[102,405]
200	83	[19,163]	-462	[-1481,-309]	285	[115,409]

Table 3: Average and CVaR of the difference: Revenue MaxMin Ratio - Revenue MinMax Regret

We should first note that the scale of the numbers in Table 3 is much smaller than in the previous paired tests, as we would expect the two regret models to share more common traits than the other models we compared with the MaxMin Ratio.

The average difference between the MaxMin Ratio and the MinMax Regret is initially negative, for small sample sizes. It later becomes positive for large samples and we notice the 90% confidence interval for this average difference is strictly positive at $N = 200$. In other words, *for large samples the MaxMin Ratio performs better on average revenues than the MinMax Regret*. On the other hand, the 5% CVaR is approximately twice as large (in absolute value) as the 95% CVaR, although without the statistical validation from the confidence intervals. *This experiment suggests the MinMax Regret could be the more robust model between the two regret-based models*.

Until this section, 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. In the next section, we diverge from this assumption and explore the use of artificial distributions for the random sampling.

3.3. Using random scenario sampling

In this experiment, we observe the behavior of the scenario based models when the sampling distribution is not the same as the true underlying distribution. In contrast with the previous section, in this experiment we will use a truncated normal distribution with a biased average as the true underlying distribution. More specifically, as the uncertainty set is defined with a box-shaped set $\delta_t \in [-15, 15]$, the underlying true distribution is sampled from a normal with mean 5 and standard deviation 7.5 and truncated over the range $[-15, 15]$. On the other hand, we will assume this is not known to the firm and we will be sampling scenarios from a uniform distribution over the uncertainty set. These uniformly sampled points will then be used to run our sampling based pricing models.

N	Average		5% CVaR		95% CVaR	
	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-165	[-3753,1649]	-3620	[-15124,-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 4: Using artificial sampling distribution: Average and CVaR of the difference Revenue SAA - Revenue MaxMin Ratio

Table 4 displays the paired comparison of the SAA revenues minus the MaxMin Ratio revenues, in a similar fashion as the comparisons performed in Section 3.2.1. *Note that the average of the difference in revenues favors the negative side, i.e. the side where MaxMin Ratio solution outperforms the SAA, suggesting that this robust solution will perform better than the SAA even on the average.* This result, although counter intuitive, appears consistent for both small and large sample sizes. The confidence intervals do not confirm this result as statistically significant at 90% confidence level, but these intervals appear to move away from the positive side as the sample size increases. The CVaR comparisons, on the other hand, are inconclusive.

N	Average		5% CVaR		95% CVaR	
	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	875	[-417,2704]	-1326	[-4438,-227]	4064	[1156,10306]
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]

Table 5: Using artificial sampling distribution: Average and CVaR of the difference Revenue MaxMin Ratio - Revenue MaxMin

Table 5 displays the results of the paired comparison of the revenues in the MaxMin Ratio model minus the MaxMin model. We can clearly see that in this case, where we used an artificial sampling distribution, the MaxMin Ratio performs both better on average and in the worst-cases. 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. *Note that the confidence intervals around these values do not overlap, strongly suggesting that the*

MaxMin Ratio is more robust than the MaxMin.

We should mention that the results with the MinMax Regret model led to the similar outcomes as the MaxMin Ratio, as expected given the similar nature of these 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 displayed here as they do not add any further insight from the previous sections.

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

3.4. Summary of results from numerical experiments

The key insights of our numerical experiments can be summarized as following:

- The SAA has a risk-neutral perspective, therefore maximizing average revenues, regardless of the final revenue distribution. It gets the best average revenue performance if the number of samples is large enough and if they are reliably coming from the true demand distribution. On the other hand, the revenue outcome is subject to the largest amount of variation, with a large shortfall in the bad revenue cases when compared to the robust models.
- The MinMax Regret model presented the more robust revenue performance when compared with other models. The MaxMin Ratio model strikes a balance between the conservativeness of the robust MinMax Regret and the aggressiveness of the SAA.
- All the Robust models, namely the MaxMin, MinMax Regret and MaxMin Ratio, will perform just as well as the SAA in average revenues for small sample sizes. For large sample sizes, the robust models will have smaller average revenues, but also with smaller shortfalls for the extreme revenue cases, in other words good for risk-averse firms. 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 worst-case revenue performance and in average revenues.
- The MaxMin Ratio and MinMax Regret have the most consistent performance when the data set provided does not come from the true underlying distribution of the demand uncertainty.

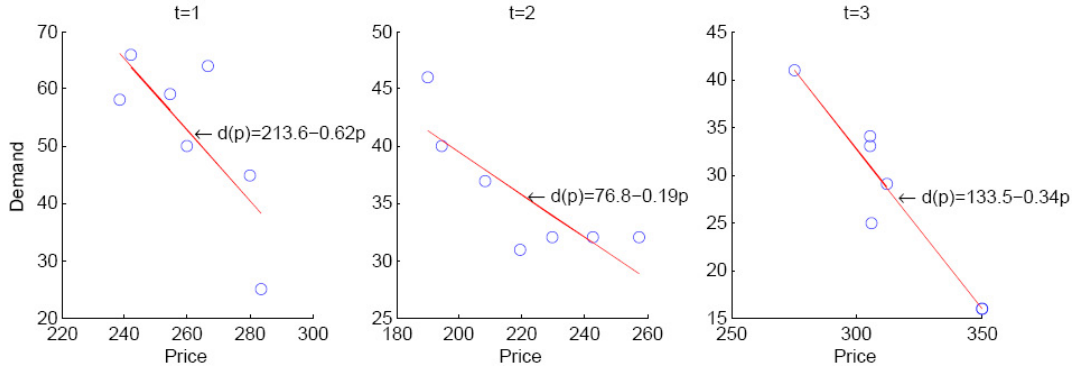


Figure 5: Price x Demand Estimation

4. Testing the model with real airline data

In this section we perform a study using real airline data, collected from a flight booking website. Tracking a specific flight with capacity of 128 seats over a period of two months for 7 different departure dates, we acquired the posted selling price and the inventory remaining for that given flight. We assumed there is no overbooking allowed and salvage value is zero, since we don't have any information about these parameters. Note that for each departure date, we have the same product being sold over and over again. To make the data more consistent, we aggregated this data into 3 periods of time: 6 or more weeks prior to departure; 5-3 weeks; and 1-2 weeks. For each of these 3 periods, we averaged the observed price and derived from the inventory data the demand observed by the company during the respective period. To summarize, we obtained 7 observations of price and demand for a 3 period horizon. Figure 5 displays the data points observed and a possible model estimation using a linear demand model. By multiplying price times demand for each period in each flight we obtain an estimate for the flight's revenue of $\$29,812 \pm 3,791$ (notation for Average \pm Standard Deviation). Note here that since we have few data points, it does not make sense to talk about the CVaR measure that was previously used to measure robustness in Sections 3.2 and 3.3. Instead, the concept of sample standard deviation will be used to express volatility of the revenue outcomes and therefore some form of robustness.

Using these 7 sample points, we perform a leave-one-out procedure to test the performance of our pricing models, where for each iteration we select one of the samples to be left out for testing and the remaining 6 points are used for training. During the training phase, a linear demand-price relation is estimated for the 6 training points and the residuals will be the deviation samples $\delta^{(i)}$, used to find the respective pricing strategies in our sample-based models. Using the test point that was left out, we obtain the difference between the point and the line estimated with the 6 training points, defining the out-of-sample deviation. It is important to emphasize that the estimation of the demand curve needs to be redone every time we leave a point out of the sample in order to obtain the independence between the testing sample and the

Model	Average	Standard Deviation
Actual Revenue	29,812	3,791
MaxMin	31,774	1,740
MinMax Regret	31,891	1,721
MaxMin Ratio	31,906	1,713
Sample Average Approximation	31,848	1,984

Table 6: Airline case: Revenue performance for each pricing model

estimated demand function. After solving the pricing models for each demand function, we can measure the out-of-sample revenue for each pricing strategy using the left-out demand point. By repeating this procedure 7 times, we obtain 7 out-of-sample revenue results, which are summarized in Table 6.

The MaxMin Ratio outperformed all other models, demonstrating the highest out-of-sample average revenue and lowest standard deviation. Compared to the actual revenue of the airline, assuming there are no external effects, the MaxMin Ratio obtained a 7% increase in the average revenue and a 55% decrease in standard deviation.

This study demonstrates, in a small scale problem, the procedure for applying our models on a real-world situation. On the other hand, there are some aspects of the airline industry that were not captured in this data set and are not considered by our models. When airlines make their pricing decisions, they take into consideration, for example, network effects between different flights and competitors' prices, which should be incorporated in future extensions of this model. A few other facts that we overlook is the aggregation done over the periods of time, due to the lack of data, or the fact that not every seat is sold at the same price at any given time, given that some people choose to buy their tickets with or without restrictions and connections. Therefore, these price-demand points are not an exact depiction of customers' response to price, but it is the best approximation we can achieve given the information that was collected. For these reasons, we cannot claim a 7% actual increase for the revenue of this flight, but instead this study suggests that there may be room for improvement in the current pricing techniques used by the airline. For instance, the methodology proposed here can be an easy improvement over the deterministic models that are used to approximate revenue-to-go functions in revenue management heuristics.

5. Conclusions

In this paper, we developed a framework to solve the dynamic pricing problem using a sampling based approach to find closed-loop robust pricing policies. Our proposed model is a practical decision tool, that can be easily implemented and solved as a convex optimization model.

Although the original robust problem is not tractable, we show that the sampled solution converges to

the robust solution as we increase the number of data points. We also provide a performance guarantee for the sampled solution based on the number of samples used. Specifically, we introduced the notion of ϵ -robustness and found the sample size needed to achieve an ϵ -robust solution with some confidence level. Moreover, we were able to extend this result from the data-driven optimization framework to a random scenario sampling approach using a relatively mild assumption on the underlying distribution.

Notice how the random scenario sampling bound is a powerful concept, given that we are now able to use a data-driven methodology to solve a robust problem. Data-driven and robust optimization have been generally considered two separate fields, since they use very different initial assumptions of information about the problem. In this research, we are building a bridge between data-driven optimization and robust optimization, which has not been explored very much. We started with an intractable robust model and illustrated how to solve a close approximation to the robust solution without any previous data, only with a likelihood bound on the underlying distribution.

Besides the theoretical guarantees mentioned above, we explored numerically the revenue performance of the different models. Among the insights, we show that the regret based models deliver a more robust revenue performance than the sample-average approximation, which is appropriate for risk averse managers. The sample-average approximation has better average revenue performance when the demand data is available in large quantities and derived from the true demand distribution. On the other hand, when the amount of data available is small or needed to be sampled from an artificial sampling distribution, the regret based models tend to perform better also on average. This small sample size performance could make this framework an important tool for managers with constraints on data availability, which is definitely an area that calls for future research.

Note that the methodological framework presented in this paper is quite general and is not limited to dynamic pricing. Because of the practicality of these models, we believe these sampling based optimization models can easily be applied to many more application areas.

Appendix A. Notation.

We summarize here the notation that will be used in throughout the paper:

- C : Initial inventory.
- T : Length of the time horizon.
- δ : Vector the demand uncertainty, containing one component for each time.
- U : Uncertainty set, $\delta \in U$.
- s : Vector of pricing control variables, i.e., pricing strategy.
- S : Strategy space, independent of the demand uncertainty.
- $p_t(s, \delta)$: Price at time t for a given pricing strategy and uncertainty realization.
- $\tilde{d}_t(s, \delta)$: Realized demand at time t , given the pricing strategy and the uncertainty.
- o : Overbooking penalty.
- v : Salvage value.
- $w(s, \delta)$: Terminal value of remaining/overbooked inventory.
- $\Pi(s, \delta)$: Actual revenue given a pricing policy and realized uncertainty.
- $\Pi^*(\delta)$: Perfect hindsight revenue, given the uncertainty realization δ .
- $h^{obj}(s, \delta)$: The objective function for any given pricing strategy s , uncertainty δ .
 obj can be one of three objectives: MaxMin, Regret or Ratio
- $h^{MaxMin}(s, \delta) = \Pi(s, \delta)$
- $h^{Regret}(s, \delta) = \Pi(s, \delta) - \Pi^*(\delta)$
- $h^{Ratio}(s, \delta) = \Pi(s, \delta)/\Pi^*(\delta)$
- z : Robust objective variable.
- $\delta^{(i)}$: i^{th} sampled realization of the demand uncertainty
- N : Number of scenario samples, as in $\delta^{(1)}, \dots, \delta^{(N)}$
- x : Set of all decision variables, which is the pricing strategy s and the robust objective z .
- n_x : Dimension of the strategy space.
- $g(x, \delta)$: Short notation for the constraints in the pricing model.
- \mathcal{Q} : True underlying distribution of the uncertainty with support in U .
- \mathcal{P} : Sampling distribution, from which we obtain scenarios of the uncertainty δ .
- f : PDF of the sampling distribution.
- k : Bounded likelihood ratio factor between \mathcal{P} and \mathcal{Q} .
- $V_{\mathcal{Q}}(x)$: Probability of violation, i.e., probability under the true distribution \mathcal{Q} , that the realized uncertainty will violate some constraints given a policy x .
- ϵ : Robust level.
- β : Risk of failure, or alternatively $(1 - \beta)$ is the confidence level.

Appendix B. Illustration of the non-convexity issue

Consider MaxMin robust problem with 2 time periods, single resource, single product, without salvage value or overbooking, a linear demand function and a linear pricing policy, the revenue function $\Pi(s, \delta)$ for any fixed s is:

$$\Pi(s, \delta) = s_1(a_1 - b_1 s_1 + \delta_1) + (s_2 + s_3 \delta_1)(a_2 - b_2(s_2 + s_3 \delta_1) + \delta_2)$$

The eigenvalues of the Hessian of $\Pi(s, \delta)$ as a function of δ , are $(-b_2 s_3 + \sqrt{b_2^2 s_3^2 + 1})s_3, (-b_2 s_3 - \sqrt{b_2^2 s_3^2 + 1})s_3$. This matrix always has one positive and one negative eigenvalue, unless the adjustable price component is zero: $s_3 = 0$ (open-loop policy). Therefore, the function is neither convex nor concave in δ .

Appendix C. Proof of Proposition 1

In this section we prove there exists an optimal solution to (3). We begin by simplifying the problem by removing the auxiliary variable z and replacing the optimization objective with $h^{obj}(s, \delta)$. It is easy to see these problems are equivalent. Note that the function $h^{obj}(s, \delta)$ is continuous with respect to s and δ from the initial assumption on the demand function, a property that is not modified by introducing the deviation component to the function. We first need to show that $\bar{h}(s) = \min_{\delta \in U} h^{obj}(s, \delta)$ is also a continuous function of s . For some $\Delta > 0$:

$$\bar{h}(s + \Delta) - \bar{h}(s) = \min_{\delta \in U} h^{obj}(s + \Delta, \delta) - \min_{\delta \in U} h^{obj}(s, \delta)$$

Define the deviations δ_1, δ_2 the deviations that attain the minimum at the two optimization problems above, $\delta_1 = \operatorname{argmin}_{\delta \in U} h^{obj}(s + \Delta, \delta)$, $\delta_2 = \operatorname{argmin}_{\delta \in U} h^{obj}(s, \delta)$. These give us upper and lower bounds:

$$h^{obj}(s + \Delta, \delta_1) - h^{obj}(s, \delta_1) \leq \bar{h}(s + \Delta) - \bar{h}(s) \leq h^{obj}(s + \Delta, \delta_2) - h^{obj}(s, \delta_2)$$

By letting $\Delta \rightarrow 0$, both sides will go to 0. It follows that: $\lim_{\Delta \rightarrow 0} |\bar{h}(s + \Delta) - \bar{h}(s)| = 0$, which implies that $\bar{h}(s)$ is a continuous function. The same argument can be used for the price function $p(s, \delta)$. We already assumed that S , the domain of s , is a compact set. We use Weierstrass Theorem to guarantee existence of solution to (3). Equivalently, we also prove the existence of solution of problem (7).

Appendix D. Proof of Theorem 2

Define x^N as the solution of the scenario based problem for a sample of size N . The probability of violation under the true uncertainty measure \mathcal{Q} is defined as $V_{\mathcal{Q}}(x^N) \doteq \mathbf{P}_{\mathcal{Q}}\{\delta \in \Delta : g(x^N, \epsilon) > 0\}$.

Define B as the set of “bad” sampling outcomes that will cause the probability of violation of the sampled solution to be greater than the ϵ robust level: $B \doteq \{(\delta^{(1)}, \dots, \delta^{(N)}) : V_{\mathcal{Q}}(x^N) > \epsilon\}$. The goal of this proof is to bound the size of B under the sampling measure \mathcal{P} , *i.e.*, $\mathbf{P}_{\mathcal{P}}(B)$.

Let n_x be the dimension of the strategy space and let I be a subset of indexes of size n_x from the full set of scenario indexes $\{1, \dots, N\}$. Let \mathcal{I} be the set of all possible choices for I , which contains $\binom{N}{n_x}$ possible subsets. Let x^I be the optimal solution to the model using only those I scenarios. Define $B_I \doteq \{(\delta^{(1)}, \dots, \delta^{(N)}) : V_{\mathcal{Q}}(x^I) > \epsilon\}$. Also based on this subset I , define the set of all possible sampling outcomes that will result in the same optimal solution as the solution using only the I scenarios: $\Delta_I^N \doteq \{(\delta^{(1)}, \dots, \delta^{(N)}) : x^I = x^N\}$.

From Calafiore and Campi (2006), we have that $B = \bigcup_{I \in \mathcal{I}} (B_I \cap \Delta_I^N)$. Therefore we need to bound the sets $B_I \cap \Delta_I^N$ under the measure \mathcal{P} .

Note that using the Bounded Likelihood Ratio assumption, we have that for each point in B , the condition $V_{\mathcal{Q}}(x^N) > \epsilon$ implies $V_{\mathcal{P}}(x^N) > \epsilon/k$. Similarly for each point in B_I , we have that $V_{\mathcal{P}}(x^I) > \epsilon/k$.

Fix any set I , for example, the first indexes $I = 1, \dots, n_x$, then note that the set B_I is a cylinder with base on the first n_x constraints, since the constraint $V_{\mathcal{Q}}(x^I) > \epsilon$ only depends on the first n_x samples. Now take any point $(\delta^{(1)}, \dots, \delta^{(n_x)}, \delta^{(n_x+1)}, \dots, \delta^{(N)})$ in the base of B_I , since we know it will satisfy $V_{\mathcal{P}}(x^I) > \epsilon/k$. For this point to be in $B_I \cap \Delta_I^N$, we only need to make sure each of the samples $\delta^{(n_x+1)}, \dots, \delta^{(N)}$ do not violate $g(x^I, \delta) \leq 0$, otherwise $x^I = x^N$ will not hold as needed in Δ_I^N . Since we have that each $N - n_x$ samples are taken independently according the sampling distribution \mathcal{P} and $V_{\mathcal{P}}(x^I) > \epsilon/k$, we have that:

$$\mathbf{P}_{\mathcal{P}} \left(B_I \cap \Delta_I^N \right) < (1 - \epsilon/k)^{N-n_x} \mathbf{P}_{\mathcal{P}}(\delta^{(1)}, \dots, \delta^{(n_x)} \in \text{base of } B_I) \leq (1 - \epsilon/k)^{N-n_x}$$

A bound on the whole set B can then be found by summing that bound over all possible subsets in \mathcal{I} .

$$\mathbf{P}_{\mathcal{P}}(B) < \binom{N}{n_x} (1 - \epsilon/k)^{N-n_x} \doteq \beta(N, \epsilon)$$

□

References

- Adida, E., Perakis, G., 2005. A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematical Programming Special Issue on Robust Optimization*.
- Ball, M. O., Queyranne, M., 2009. Toward robust revenue management: Competitive analysis of online booking. *Operations Research* 57 (4), 950–963.
- Ben-Tal, A., Golany, B., Nemirovski, A., Vial, J., 2005. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing and Service Operations Management* 7 (3), 248–271.
- Ben-Tal, A., Goryashko, A., Guslitzer, E., Nemirovski, A., 2004. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming* 99 (2), 351–376.
- Ben-Tal, A., Nemirovski, A., 1998. Robust convex optimization. *Mathematics of Operations Research* 23, 769–805.
- Ben-Tal, A., Nemirovski, A., 1999. Robust solutions to uncertain linear programs. *Operations Research Letters* 25, 1–13.
- Ben-Tal, A., Nemirovski, A., 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming* 88, 411–424.

- Bertsimas, D., Iancu, D., Parrilo, P., 2010. Optimality of affine policies in multi-stage robust optimization. *Mathematics of Operations Research* 35 (2), 363–394.
- Bertsimas, D., Sim, M., 2004. The price of robustness. *Operations Research* 52 (1), 35–53.
- Bertsimas, D., Thiele, A., 2006. A robust optimization approach to supply chain management. *Operations Research* 54 (1).
- Birbil, S. I., Frenk, J. B. G., Gromicho, J. A. S., Zhang, S., 2006. An integrated approach to single-leg airline revenue management: The role of robust optimization. *Econometric Institute Report E.I. 2006-20*.
- Bitran, G., Caldentey, R., 2003. An overview of pricing models and revenue management. *Manufacturing and Service Operations Management* 5 (3), 203–229.
- Calafiore, G., November 2009. On the expected probability of constraint violation in sampled convex programs. *Journal of Optimization Theory and Applications* 143 (2), 405–412.
- Calafiore, G., Campi, M., 2005. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming* 102, 25–46.
- Calafiore, G., Campi, M., May 2006. The scenario approach to robust control design. *IEEE Transactions On Automatic Control* 51 (5), 742–753.
- Campi, M., Garatti, S., 2008. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization* 19 (3), 1211–1230.
- Caramanis, C., May 2006. Adaptable optimization: Theory and algorithms. PhD dissertation, Department of Electrical Engineering and Computer Science, MIT.
- Elmaghraby, W., Keskinocak, P., 2003. Dynamic pricing: Research overview, current practices and future directions. *Management Science* 49, 1287–1309.
- Eren, S. S., Maglaras, C., 2010. Monopoly pricing with limited demand information. *Journal of Revenue and Pricing Management* 9, 23–48.
- Gallego, G., Moon, I., 1993. The distribution free newsboy problem: Review and extensions. *The Journal of the Operational Research Society* 44 (8).
- Gaur, V., Kesavan, S., Raman, A., Fisher, M., 2007. Estimating demand uncertainty using judgmental forecasts. *Manufacturing & Service Operations Management* 9 (4), 480–491.
- Goldfarb, D., Iyengar, G., 2003. Robust convex quadratically constrained programs. *Mathematical Programming* 97 (3), 495–515.
- Kleywegt, A., Shapiro, A., Homem-de Mello, T., 2001. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12 (2), 479–502.
- Lan, Y., Gao, H., Ball, M., Karaesmen, I., 2008. Revenue management with limited demand information. *Management Science* 54 (9).
- Levi, R., Perakis, G., Uichanco, J., 2012. The data-driven newsvendor problem: New bounds and insights. *Accepted in Operations Research*.
- Levi, R., Roundy, R. O., Shmoys, D. B., 2007. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research* 32 (4), 821–839.
- Lim, A., Shanthikumar, J., 2007. Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research* 55 (2), 198–214.
- Lim, A., Shanthikumar, J., Watwai, T., 2008. Robust multi-product pricing and a new class of risk-sensitive control problems. Working paper, Department of Industrial Engineering and Operations Research, University of California, Berkeley, CA.
- Lobel, R., May 2012. Pricing and incentive design in applications of green technology subsidies and revenue management. PhD dissertation, Operations Research Center, MIT.

- Nguyen, D. T., Perakis, G., 2005. Robust competitive pricing for multiple perishable products, working paper, MIT.
- Pagnoncelli, B., Ahmed, S., Shapiro, A., 2009. Computational study of a chance-constrained portfolio problem. *Journal of Optimization Theory and Applications (JOTA)* 142 (2), 399–416.
- Perakis, G., Roels, G., 2008. Regret in the newsvendor model with partial information. *Operations Research* 56 (1), 188–203.
- Perakis, G., Sood, A., 2006. Competitive multi-period pricing for perishable products: A robust optimization approach. *Mathematical Programming* 107 (1).
- Roels, G., Perakis, G., 2010. Robust controls for network revenue management. *Manufacturing & Service Operations Management* 12 (1), 56–76.
- Rusmevichientong, P., Van Roy, B., Glynn, P. W., 2006. A non-parametric approach to multi-product pricing. *Operations Research* 54 (1), 89–98.
- Savage, L. J., 1951. The theory of statistical decisions. *American Statistical Association Journal* 46, 55–67.
- Scarf, H. E., 1958. A min-max solution to an inventory problem. *Studies in Mathematical Theory of Inventory and Production*, 201–209.
- Soyster, A., 1973. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations Research* 21, 1154–1157.
- Talluri, K., van Ryzin, G., 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers.
- Thiele, A., 2009. Multi-product pricing via robust optimization. *Journal of Revenue and Pricing Management* 8 (1), 67–80.
- Yue, J., Chen, B., M.-C., W., 2006. Expected value of distribution information for the newsvendor problem. *Operations Research* 54 (6), 1128–1136.
- Zhan, R., Shen, Z., 2005. Newsvendor problem with pricing: properties, algorithms, and simulation. In: *WSC '05: Proceedings of the 37th conference on Winter simulation. Winter Simulation Conference*, pp. 1743–1748.
- Zhang, L., Sep. 2006. Multi-period pricing for perishable products: Uncertainty and competition. MS dissertation, *Computation Design and Optimization*, MIT.