

MIT Open Access Articles

Dynamic Pricing through Sampling Based Optimization

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

Citation: Lobel, Ruben, and Georgia Perakis. "Dynamic pricing through sampling based optimization." The 51st Airline Group of the International Federation of Operational Research Societies (AGIFORS) Annual Proceedings, Antalya, Turkey, October 10-14, 2011.

As Published: <http://toc.proceedings.com/13874webtoc.pdf>

Publisher: Airline Group of the International Federation of Operational Research Societies

Persistent URL: <http://hdl.handle.net/1721.1/88129>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

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



Dynamic Pricing Through Sampling Based Optimization

Ruben Lobel

Operations Research Center, MIT, Cambridge, MA 02139, rlobel@mit.edu

Georgia Perakis

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

In this paper we develop an approach to dynamic pricing that combines ideas from data-driven and robust optimization to address the uncertain and dynamic aspects of the problem. In our setting, a firm offers multiple products to be sold over a fixed discrete time horizon. Each product sold consumes one or more resources, possibly sharing the same resources among different products. The firm is given a fixed initial inventory of these resources and cannot replenish this inventory during the selling season. We assume there is uncertainty about the demand seen by the firm for each product and seek to determine a robust and dynamic pricing strategy that maximizes revenue over the time horizon. While the traditional robust optimization models are tractable, they give rise to static policies and are often too conservative. The main contribution of this paper is the exploration of closed-loop pricing policies for different robust objectives, such as MaxMin, MinMax Regret and MaxMin Ratio. We introduce a sampling based optimization approach that can solve this problem in a tractable way, with a confidence level and a robustness level based on the number of samples used. We will show how this methodology can be used for data-driven pricing or adapted for a random sampling optimization approach when limited information is known about the demand uncertainty. Finally, we compare the revenue performance of the different models using numerical simulations, exploring the behavior of each model under different sample sizes and sampling distributions.

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

History:

1. Introduction

In many industries, managers are faced with the challenge of selling a fixed amount of inventory within a specific time horizon. Moreover, the firm offers an array of different products which will consume a shared set of resources. 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 or fail to capture potential revenues. In the pricing problem, demand is usually modeled as a function of the prices set by the firm. To model uncertainty in demand, we represent some parameters of this function as random. These random parameters can be modeled assuming a fixed distribution, as in the stochastic optimization literature, or using the distribution-free framework of robust optimization. In the stochastic optimization approach, demand uncertainty is given by a certain probability distribution and the firm's objective is to find prices that maximize its expected revenue. A disadvantage of this approach is that it requires full knowledge of the demand distribution, often requiring the uncertainty to be independent across time periods. Note that the dynamic programming formulation used to solve the stochastic optimization model will have a serious dimensionality problem if there are many types of resources or if demands are correlated. On the other hand, in the robust optimization approach, 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 demands 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 uncertainty 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.

In order to achieve this goal, we are going to develop pricing optimization models that are based on sampled scenarios of demand. At first, using the assumption that we have access to historical samples of demand, we will develop a data-driven robust approach to pricing. Furthermore, in the case these samples are not available, we will develop a framework to generate random scenario samples using mild assumptions on the demand distribution and still obtain a sampling based robust solution. We will also provide theoretical bounds on the number of samples required to obtain a given confidence level about our solutions and develop numerical examples to demonstrate how to apply this methodology in a practical setting and to obtain managerial insights about these pricing models.

1.1. 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). The revenue management literature in general assumes a fixed set of pre-determined prices and tries to allocate capacity for different price levels. This problem can be traced back to Littlewood (1972), who first tackled the network capacity control problem. The complexity of solving these large network models led to the development of numerous heuristics. Belobaba (1987, 1992) later developed heuristics such as EMSR-a and EMSR-b which performed very well in practice for solving the capacity control problem. Talluri and van Ryzin (1998) further developed the concept of bid-price controls. Bertsimas and Popescu (2003) solved this problem by approximating the revenue-to-go function and Bertsimas and de Boer (2005) introduced an algorithm based on stochastic gradient. Akan and Ata (2009) recently developed a martingale characterization of bid-price controls. On the other hand, the dynamic pricing literature attempts to solve a similar problem, but does not assume fixed prices, rather it assumes that we can choose different prices at each period. In that case, we also assume that the aggregate demand seen by the firm is a function of the price.

Most of the revenue management and pricing literature uses the stochastic optimization framework, which makes 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, we use past historical data, which are sample evaluations coming from the true demand distribution. In revenue management, this approach has been pioneered by van Ryzin and McGill (2000) who introduced an adaptive algorithm for booking limits. Bertsimas and de Boer (2005) and van Ryzin and Vulcano (2008) developed a stochastic gradient algorithm to solve a revenue management problem using the scenario samples. Another 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 and Levi et al. (2010), 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 (2009). Besbes and Zeevi (2008) assume there is no prior data and develop an approach that requires an exploration phase to obtain data and an exploitation phase to generate revenue. Other relevant work on the subject of revenue management with demand learning is by Araman and Caldentey (2009) and by Farias and Van Roy (2009).

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 (2006), 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 (2007). An alternative approach is the relative

regret measure, also known as the competitive ratio. In revenue management and pricing, Ball and Queyranne (2006) and Lan et al. (2006) 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 (2006). 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. (2009) 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.

As for the numerical experiments, we will compare the out-of-sample revenues of our sampling-based pricing models by using Monte Carlo-based performance measures. Chiralaksanakul and Morton (2004) and Bayraksan and Morton (2006) proposed Monte Carlo-based procedures to assess the quality of stochastic programs, building sampling-based confidence intervals for the optimality gap of the solution. To compare our pricing models, we will conduct paired tests and use similar

concepts of sampling based confidence intervals around the average and worst-case difference in revenues.

1.2. Contributions

The main goal of this paper is to solve the multi-item dynamic pricing problem in a practical way. In particular, we develop a methodology that is easy to implement, uses directly available data in the price optimization model, and at the same time is justified from a theoretical point of view. We accomplish this by using a sampling based optimization approach, which allows for a wide range of modeling complexity such as adjustable pricing policies, nonlinear demand functions, network effects, overbooking and salvage value. The solution concept that motivates our model is the robust pricing problem. Nevertheless this problem becomes intractable as we introduce all the modeling components mentioned before. We show in this paper how this new combined robust sampling based framework is tractable, performs very well in numerical experiments and is theoretically grounded as a good approximation to the original robust problem, depending on the number of samples used.

More specifically, most robust approaches to the pricing problem use an open-loop model, where the information obtained by observing past demand is not used in the pricing decision of future periods. A key challenge, when introducing adjustability to dynamic pricing, is that the model easily turns into a non-convex problem, which is intractable to solve even when using traditional robust optimization techniques. Instead we propose a sampling based approach where we solve the problem using only a given set of demand scenarios. The new problem becomes a convex program and can be efficiently solved. The 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. Nevertheless, this type of data-driven approach uses the assumption that we have 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. We are able to show bounds on the sample size required for the random sampling solution using a relatively mild assumption on the unknown underlying distribution, which we refer to as a bounded likelihood ratio assumption.

Although fairly general, the bounded likelihood ratio assumption can be seen as too abstract and hard to quantify in a practical setting. For this reason we further developed a method to obtain this likelihood ratio measure using more commonly used assumptions, which can be more easily verified by the firm. More specifically, using the assumption that the uncertain parameter of the demand follows a log-concave distribution, independent across products and time periods and with a known standard deviation, we obtain the likelihood ratio bound on the true unknown distribution, relative to a uniform random sampling distribution. Therefore we can use the bound developed before to determine the sample size that needs to be generated to obtain any level of robustness for the pricing problem.

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. We start with an intractable robust model and illustrate how to solve a close approximation to the robust solution using only randomly generated scenarios.

Another contribution of this work comes from the numerical experiments, where we study the simulated revenue performance of our dynamic pricing models. Our first experiment will provide empirical evidence that the sample size bound we present in this paper is in fact tight for our dynamic pricing problem. In the other experiments, we will compare the performance of the robust pricing models, using three different types of robust objectives (MaxMin, MinMax Regret, MaxMin Ratio) and a Sample Average Approximation (SAA) model, which is more common in the stochastic optimization literature. When compared to the SAA benchmark, the robust models will usually obtain a smaller average revenue but will do significantly better in the lower revenue cases. On the other hand, when the number of samples provided are small, the robust models have shown to perform just as well as the SAA on average, while still maintaining a less volatile revenue outcome. Comparing between the robust models, we show that the traditional robust (MaxMin) is usually dominated by the regret based models, both in average and worst-case revenue. The two regret-based models (MinMax Regret and MaxMin Ratio) have a relatively similar performance, but the first one tends to be more robust and conservative than the second.

We also show in this paper how to apply our methodology in practice, using a case study with actual airline data. Without considering competition effects, we will show how our robust pricing models might achieve significant improvements in the revenues per flight.

The remainder of the paper is structured as follows. In Section 2, we introduce the modeling approach that we will consider and discuss the implementation issues. In Section 3, we show the simulated performance of the proposed models and interpret the numerical results. In Section 4, we develop an airline case study. Finally, in Section 5, we conclude with a summary of the discussions mentioned before and give possible directions for future work.

The online 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-F display the proofs of our theoretical results in Section 2. Appendices G-J provide information about the numerical experiments displayed Section 3, as well as additional experiments.

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 $\delta_t^{(1)}, \dots, \delta_t^{(N)}$ 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:

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

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 network effects, salvage/overbooking inventory and nonlinear demand models. Throughout this section we develop a solution approach that mitigates the modeling issues described before. Further on, in Section 2.1, 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.

To address the network component of pricing problems, we consider the firm to have n products (eg. itineraries) and m resources (eg. flight legs). The $m \times n$ incidence matrix M determines which resources are consumed by each product sold. Also let T be the length of the time horizon. Define $p_{j,t}$ as the price of product j at time t and the demand function as $d_{j,t}$. We admit that the demand $d_{j,t}$ can be a function of the prices of all the products at any time period. For the sake of clarity, we will use demand functions that are affected only by the current product price $d_{j,t}(p_{j,t})$. This means that a certain product's demand will not be affected by the prices of the other products or by the past prices. It is important to note that the modeling techniques and the results presented in this paper can be easily implemented with cross product/time effects (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. The firm's goal is to set a pricing policy for each product that robustly maximizes the revenue of the firm, $\sum_{t=1}^T \sum_{j=1}^n p_{j,t} d_{j,t}(p_{j,t})$. For technical reasons, which are discussed later, we require that the demand function satisfies the following convexity/concavity assumption.

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

For example, suppose we assume a linear demand function, $d_{j,t}(p_{j,t}) = a_{j,t} - b_{j,t}p_{j,t}$, where $a_{j,t}$ and $b_{j,t}$ are scalars. 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, besides the linear demand function, which are common in the revenue management literature are the iso-elastic

demand $d_{j,t}(p_{j,t}) = a(p_{j,t})^{-b}$ for $b \in [0, 1]$, and the logarithmic demand $d_{j,t}(p_{j,t}) = -a \log(p_{j,t}/b)$. In the numerical experiment of Appendix H, we will illustrate our framework using the logarithmic demand model.

Without considering demand uncertainty, the objective is to find the optimal set of prices that maximizes the sum of the revenues for the entire horizon. The prices are nonnegative and the total demand seen by the firm for a given resource l should be less than its total capacity C_l or else the firm will pay an overbooking fee o_l for every unit of resource sold above capacity of resource l . For every unit of capacity not sold, the firm will get a salvage value of g_l . We require that $o_l > g_l$ for all resources $l = 1, \dots, m$ 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 very easy to justify. Define w_l as the difference between the number of units sold and the capacity in each resource l , while $w_l^+ = \max(w_l, 0)$ and $w_l^- = \min(w_l, 0)$.

To capture the variability in demand, given the nominal demand function $d_{j,t}(p_{j,t})$, define $\tilde{d}_{j,t}(p_{j,t}, \delta_{j,t})$ as the actual demand for product j at time t , which is realized for some uncertain parameter $\delta_{j,t}$. For example, suppose we have a linear nominal demand function given by $d_{j,t} = a_{j,t} - b_{j,t}p_{j,t}$. Then introducing an additive uncertainty we would have: $\tilde{d}_{j,t}(p_{j,t}, \delta_{j,t}) = a_{j,t} - b_{j,t}p_{j,t} + \delta_{j,t}$. 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.

In general, $\delta = (\delta_{1,1}, \dots, \delta_{n,T})$ is a random vector with one component $\delta_{j,t}$ for each product j and each period t (it can be easily generalized for multiple uncertain components in each product or time period). 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 or products, as opposed to most stochastic optimization approaches. Although not required in our general framework, we will use an independence assumption when giving an example of how to apply the random sampling approach (see Corollary 1).

In our framework, we assume there is a finite set of pricing decision variables s , which lie within the strategy space S and define the prices $p_{j,t}(s, \delta)$ for each product at each period. 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 $p_{j,t}(s, \delta)$ for both cases, while $p_{j,t}(s, \delta)$ is actually independent of δ in open-loop policies and in closed-loop policies a function of $\delta_{j,1}, \dots, \delta_{j,t-1}$, including all products j . Also to avoid redundancy

in notation, define $\tilde{d}_{j,t}(s, \delta) = \tilde{d}_{j,t}(p_{j,t}(s, \delta), \delta_{j,t})$. In other words, the policy s and the uncertainty δ determine the prices of all products at time t , therefore also determining the realized demand. Define the net leftover inventory of resource l as a function of the pricing policy and the realized uncertainty: $\tilde{w}_l(s, \delta) = \sum_{t=1}^T \sum_{j=1}^n M_{l,j} \tilde{d}_{j,t}(s, \delta) - C_l$.

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**.

$$\begin{aligned}
h^{MaxMin}(s, \delta) &= \sum_{t=1}^T \sum_{j=1}^n p_{j,t}(s, \delta) \tilde{d}_{j,t}(s, \delta) - \sum_{l=1}^m o_l \tilde{w}_l(s, \delta)^+ + g_l \tilde{w}_l(s, \delta)^- \\
h^{Regret}(s, \delta) &= \sum_{t=1}^T \sum_{j=1}^n p_{j,t}(s, \delta) \tilde{d}_{j,t}(s, \delta) - \sum_{l=1}^m o_l \tilde{w}_l(s, \delta)^+ + g_l \tilde{w}_l(s, \delta)^- \\
&\quad - \max_{y \geq 0} \left\{ \sum_{t=1}^T \sum_{j=1}^n y_{j,t} \tilde{d}_{j,t}(y, \delta_t) - \sum_{l=1}^m o_l \tilde{w}_l(y, \delta)^+ + g_l \tilde{w}_l(y, \delta)^- \right\} \\
h^{Ratio}(s, \delta) &= \frac{\sum_{t=1}^T \sum_{j=1}^n p_{j,t}(s, \delta) \tilde{d}_{j,t}(s, \delta) - \sum_{l=1}^m o_l \tilde{w}_l(s, \delta)^+ + g_l \tilde{w}_l(s, \delta)^-}{\max_{y \geq 0} \left\{ \sum_{t=1}^T \sum_{j=1}^n y_{j,t} \tilde{d}_{j,t}(y, \delta_t) - \sum_{l=1}^m o_l \tilde{w}_l(y, \delta)^+ + g_l \tilde{w}_l(y, \delta)^- \right\}}
\end{aligned}$$

In the first robust concept, **MaxMin**, we try to find a pricing policy that gives us the best revenue for all possible realizations of the uncertainty δ in the set U . More specifically, the actual revenue realized for some policy s and deviation δ is given by $\Pi(s, \delta) = \sum_{t=1}^T \sum_{j=1}^n p_{j,t}(s, \delta) \tilde{d}_{j,t}(s, \delta) - \sum_{l=1}^m o_l \tilde{w}_l(s, \delta)^+ + g_l \tilde{w}_l(s, \delta)^-$. Then for the MaxMin case, the objective function is simply given by $h^{MaxMin}(s, \delta) = \Pi(s, \delta)$. 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**, also known as absolute regret. In this case, 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:

$$\begin{aligned}
\Pi^*(\delta) = \max_{y, w} & \sum_{t=1}^T \sum_{j=1}^n y_{j,t} \tilde{d}_{j,t}(y, \delta) - \sum_{l=1}^m o_l w_l^+ + g_l w_l^- \\
s.t. & \sum_{t=1}^T \sum_{j=1}^n M_{l,j} \tilde{d}_{j,t}(y, \delta) \leq C_l + w_l, \quad \forall l = 1, \dots, m \\
& y \geq 0
\end{aligned} \tag{1}$$

The model above for the hindsight revenue $\Pi^*(\delta)$ is a deterministic convex optimization problem, which can be efficiently computed for any given δ (either in closed form or using an optimization package). Note that the capacity constraint can always be satisfied with equality in the optimal solution, since left-over capacity can be shifted to w_l while improving the objective function because of the salvage value, $g_l \geq 0$. Therefore, at optimality, the left-over capacity variable from the hindsight model above will have $w_l = \tilde{w}_l(y, \delta)$.

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 $h^{Ratio}(s, \delta) = \frac{\Pi(s, \delta)}{\Pi^*(\delta)}$.

Consider the following model which we call the robust pricing model. We will go into details of the different modeling techniques later in this section, but they all fall within the following general framework:

$$\begin{aligned}
& \max_{s \in S, z} z \\
s.t. & \left\{ \begin{array}{l} z \leq h^{obj}(s, \delta) \\ p(s, \delta) \geq 0 \end{array} \right\} \forall \delta \in U
\end{aligned} \tag{2}$$

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 $\tilde{d}_t(p_t, \delta_t)$ for any given price p_t and uncertainty δ_t ; 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 2.1). 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.

To introduce adjustability in our model, we must find a pricing policy that is a function of the previous realized uncertainties. On the other hand, searching over all possible functions is a rather hard task. An approach used in the literature to introduce adaptability to the robust model is the affinely adjustable robust optimization, where the search is limited to policies that are linear on

the control variables. For example, define the price p using the variables u as the deterministic component of the price and v as the effect that past demand uncertainty should have on current prices. In this case $s = (u, v)$. Then a typical pricing policy can be defined as:

$$p_{j,t}(u, v, \delta) = u_{j,t} + \sum_{l=1}^m v_{j,l,t} \sum_{j'=1}^n \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'} \quad (3)$$

This policy (3) will be used in our numerical experiments in Section 3. In this policy example above, we have an adjustable variable for each resource, where the summation $\sum_{j'=1}^n \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}$ contains a lot of the inventory information necessary for our adjustable policy. As an intuition for choosing this type of policy, observe that when using additive uncertainty the sum of past uncertainties $\sum_{j'=1}^n \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}$ is the amount of sales realized above or below the deterministic level expected up to time $t-1$ for the resource l . If the firm has sold more than expected, we should have an incentive to raise prices, given the reduced amount of inventory left. On the other hand, if uncertainties are negatively correlated across time, a series of low demands could induce a higher demand in the future, possibly giving the incentive to raise prices. Notice how static policies can be overly conservative because they cannot adjust the prices according to the different levels of the inventory, while the adjustable policies are more flexible. In the example above, we chose a pricing policy that is linear on δ . In some situations, we might rather capture the effect of past deviations in other forms, for instance δ^2 or $|\delta|$, depending on the structure of the demand function and the uncertainty set. In the framework we propose, pricing policies can be any function of previous realizations of δ , as long as the function is linear in the pricing decision variables for any given δ .

ASSUMPTION 2. *Let $p_t(s, \delta_1, \dots, \delta_{t-1})$ be the pricing policy at time t for a given set of pricing decision variables s and uncertainties realized up to time t . We will restrict ourselves to policies that are linear on s for any given δ . Nevertheless the dependence on δ can be nonlinear. We also assume that s is restricted by the strategy space S , which is a finite dimensional, compact and convex set.*

The reason for Assumption 2 can be easily illustrated in the pricing policy in (3). When introducing this policy as an equality in the optimization model (2), we must guarantee that it defines a convex set in p , u and v for any given δ . This convexity will hold if the pricing policy is linear on the decision variables (u, v) , as stated in Assumption 2.

Furthermore, the set S defined above captures only the constraints in the strategy space that do not depend on the uncertainty. In general, we expect S to be a box set, with upper and lower

bounds for the pricing variables. The finite dimensionality and compactness of set S are introduced to guarantee the existence of a solution to the optimization problem.

Note that using adjustable policies, the inequality $z \leq h^{obj}(p(u, v, \delta), \delta)$, 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 and it is not necessary to show that more complicated modeling will generally not resolve the convexity issue. In the next section we introduce the sampling based approach that we advocate for solving the robust pricing problem.

2.1. 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 (2) 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 user 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:

$$\begin{aligned} & \max_{s \in S, z} z \\ & s.t. \quad \left\{ \begin{array}{l} z \leq h^{obj}(s, \delta^{(i)}) \\ p(s, \delta^{(i)}) \geq 0 \end{array} \right\} \forall i = 1 \dots N \end{aligned} \quad (4)$$

The idea of using sampled uncertainty scenarios, or data points, in stochastic optimization models is often called sampling or scenario based optimization. Note that from Assumptions 1 and 2, the constraints in (4) define a convex set in (s, z) 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 (2) has an optimal solution (for a proof under more concise notation, see Appendix C). Also, from Assumption 1, it follows that for any fixed uncertainty δ , there exist a unique optimal solution to the exact robust

pricing problem. 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. To be consistent with our distribution-free approach, we should not rely on hypothesis tests that use a normality assumption about the error. If such assumption can be made about the estimation errors in the first place, then it could potentially be explored in the optimization model. In this case, we use the linear regression simply as a tool to obtain estimates of the demand parameters we need, even if we cannot validate the fit with statistical testing. 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.

Define the set of decision variables $x = (s, z)$, such that it lies within the domain $x \in X$, where $X = S \times \Re$. Define c with the same dimension as x such that $c = (0, \dots, 0, 1)$. Also define the equivalent constraint function with a scalar valued f such that: $f(x, \delta) = \max \{ z - h^{obj}(s, \delta), -p(s, \delta) \}$.

Since each constraint in the definition of $f(x, \delta)$ above is convex in x for any given δ , the maximum between them is still convex in x . Moreover, $f(x, \delta) \leq 0$ is equivalent to the constraint set defined before in (2). Then problem (2) can be concisely defined as the following model (5):

$$\max_{x \in X} c'x, \quad \text{s.t. } f(x, \delta) \leq 0, \forall \delta \in U \quad (5)$$

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 (4) can be defined as the following model (6):

$$\max_{x \in X} c'x, \quad \text{s.t. } f(x, \delta^{(i)}) \leq 0, \forall i = 1 \dots N \quad (6)$$

The following definition is required to develop the concept of ϵ -robustness which we will use throughout the paper.

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:

$$V_{\mathcal{Q}}(x) = \mathbf{P}_{\mathcal{Q}}\{\delta : f(x, \delta) > 0\}.$$

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 (6). 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. In the following Theorem 1, we show how one can derive this ratio k from the standard deviation σ , when the true uncertainty distribution belongs to the family of log-concave distributions with bounded support and is independent across products and time periods. In this case, the standard deviation of the demand, which is a very familiar statistic to most managers, should be somehow obtained by the firm. Similar theorems could 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, which includes uniform, normal, logistic, extreme-value, chi-square, chi, exponential, laplace, among others. For further reference, a deeper understanding of log-concave distributions and their properties can be seen in Bagnoli and Bergstrom (1989) and Bobkov (1999).

THEOREM 1. *Assume that the deviations $\delta_{j,t}$ for each product $j = 1, \dots, n$ and time period $t = 1, \dots, T$ are independent and bounded in a box-shaped uncertainty set: $\delta_{j,t} \in [lb, ub]$. We do not know the true distribution of $\delta_{j,t}$, but we assume it belongs in the family of log-concave distributions. Furthermore, the standard deviation of $\delta_{j,t}$ is known to be σ . We randomly sample points from a uniform distribution supported over the uncertainty set, with density $\frac{1}{ub-lb}$ for each $\delta_{j,t}$.*

Relative to the uniform, the true distribution has a likelihood ratio bound $k = \left[(ub - lb) \frac{e^{\sqrt{6}}}{\sigma\sqrt{2}} \right]^{nT}$. If we further assume symmetry for the distribution of δ , the bound will be $k = \left[(ub - lb) \frac{1}{\sigma\sqrt{2}} \right]^{nT}$.

For a proof of Theorem 1, see Appendix D. 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. *Given the following scenarios of uncertainty $\delta^{(1)}, \dots, \delta^{(N)}$, drawn from an artificial distribution \mathcal{P} , which bounds the true uncertainty distribution \mathcal{Q} by a factor of k (see Definition 3 and Theorem 1). Let n_x be the dimension of the strategy space X , x^N be the solution of (6) using the N sample points, and ϵ be the robust level parameter. Then define a “risk of failure” parameter $\beta(N, \epsilon)$ 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 found using this method 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. 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.

For proof of Theorem 2, see Appendix E. The proof of this theorem is similar to the proof of Calafiore and Campi (2006) (Theorem 1), but 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 carefully introduce the bounded likelihood ratio assumption to connect the true distribution to the sampling distribution. Corollary 1 can be derived using algebraic manipulations by solving for N the expression $\beta \leq \binom{N}{n_x} (1 - \epsilon/k)^{N-n_x}$.

Note that this result can be quite useful in practice, specially if there is a small amount of data or no data samples at all. It is common for managers to assume the demand to lie within a certain range. If they can further make an assessment of the volatility of demand (standard deviation), we can apply this result. From Theorem 1, we obtain a likelihood ratio bound k , which directly implies the number of samples required to be drawn for a given confidence and robustness level, according to Theorem 2. 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 (log-concavity, range and standard deviation).

In more recent literature, Campi and Garatti (2008) developed a better bound for the case where the samples come from the true distribution, which we call the data-driven case ($k = 1$). In fact they prove that this new bound is the tightest possible bound, since it holds with equality for the special case of fully-supported problems. They also show how the new bound will always provide a smaller sample size requirement than the one developed in Theorem 2. On the other hand, we have not been able to generalize this result for the case of random scenario sampling with artificial distributions ($k > 1$), which is a much more complicated problem. In Appendix F, we present the result from Campi and Garatti (2008) and in the subsection Appendix F.1 we demonstrate numerically that this bound is actually tight for an instance of the dynamic pricing problem. As shown in Campi and Garatti (2008), the sample size bound in Corollary 1, when $k = 1$, requires significantly more samples for the same β and ϵ levels than Corollary F.1 in the appendix.

In order to apply these confidence bounds to our pricing problem, we only need to know the likelihood bound k , which is either data provided by the firm or calculated in a similar manner as in Theorem 1, 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, the number of decision variables is the number of time periods T times the number of products n (one fixed price per product/period) plus one for the robust objective variable: $n_x = Tn + 1$. For the pricing policy suggested in (3), which we use in Section 3, $n_x = 1 + Tn + (T - 1)nm$. For instance, in the airline case study of Section 4, we deal with a 3-period single product adjustable model, therefore $n_x = 6$.

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. The experiments were made on a server using a single core of 2.8GHz Intel Xeon processor with 32GB of memory.

The first experiment, in Section 3.1, is designed to compare the out-of-sample performance of the multiple pricing models discussed in Section 2 in order to obtain managerial insights on the effects of the models on the resulting pricing policies. In Section 3.2, we observe the performance of the different models when the sampling distribution is not the same as the true underlying distribution.

More specifically, we will compare the following models: MaxMin, MinMax Regret, MaxMin Ratio, and Sample Average Approximation. A summary of all the optimization models used in this section is available in Appendix G. 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 and present the optimization model in Appendix G.

In Appendix H, we display the out-of-sample revenues between the static and the adjustable policies. This experiment illustrates why we have not dedicated much time discussing the static models in this paper, as it is clearly dominated by the adjustable counterpart. For this reason, all other experiments in this paper use adjustable policies. In Appendix I we demonstrate an application of the model using a network model with multiple products and resources and in Appendix J we demonstrate the application of a nonlinear demand model. The qualitative conclusions of these two cases (network and nonlinear demand) do not differ from the single product, linear demand models of Section 3.1 and 3.2. To avoid redundancy, we will not dedicate much time to these examples.

3.1. Using true distribution sampling: data-driven approach

In order to isolate the differences between the models, we will reduce the problem to a single product case ($n = m = 1$) and linear demand functions with additive uncertainty. The reader should note that this is used here for simplicity of notation only and the use of other demand and uncertainty models can be easily incorporated.

Define the **MaxMin** model as in the optimization problem formulated in (G.1), in Appendix G. For the regret-based models, we can solve the hindsight revenue $\Pi^*(\delta^{(i)})$ from (G.2) for each data point $\delta^{(i)}$ of the uncertainty sample by solving a simple convex optimization problem. The **MinMax Regret** is defined in (G.3) by subtracting the term $\Pi^*(\delta^{(i)})$ from the right hand side in the objective constraint (the first constraint). Similarly, the **MaxMin Ratio** is defined by dividing the right hand side by the term $\Pi^*(\delta^{(i)})$, as in the optimization problem (G.4).

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. If we define the pricing decision as a function of the previous realized uncertainty, we have an adjustable stochastic optimization problem, which is the same as the adjustable robust models above, except that the objective approximates the expected value of the revenue by taking the average of each revenue point: $\frac{1}{N} \sum_{i=1}^N \left[\sum_{t=1}^T \sum_{j=1}^n p_{j,t} (a_t - b_t p_{j,t} + \delta_t^{(i)}) - \sum_{l=1}^m o_l w_l^+ + g_l w_l^- \right]$. This solution concept is what we refer to as Sample Average Approximation (**SAA**) and we detailed the optimization model (G.5).




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, but it illustrates how the revenue outcomes can be quite different between models when sample sizes are too small. The intrinsic volatility of the 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, the intrinsic volatility of revenues is reduced by the difference operator, which allows us to compare the models’ performances.

Using the paired comparisons, we will 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 quantile α , the CVaR_α is the expected value of a random variable conditional that the variable is below its α quantile. This statistic is as a way to measure the tail of a distribution and determine how bad things can go when they go wrong. Since we are not looking at one particular revenue distribution, but instead the difference between two revenues, it is only natural to look at the positive tail of the distribution, when the second comparing model is at risk. 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 quantile.

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

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.




Figure 2 Paired 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 $\text{CVaR}_{5\%} = -2233$ and the $\text{CVaR}_{95\%} = 874$. 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 $\text{CVaR}_{5\%}$ of this difference is significantly larger than $\text{CVaR}_{95\%}$ (in absolute terms).

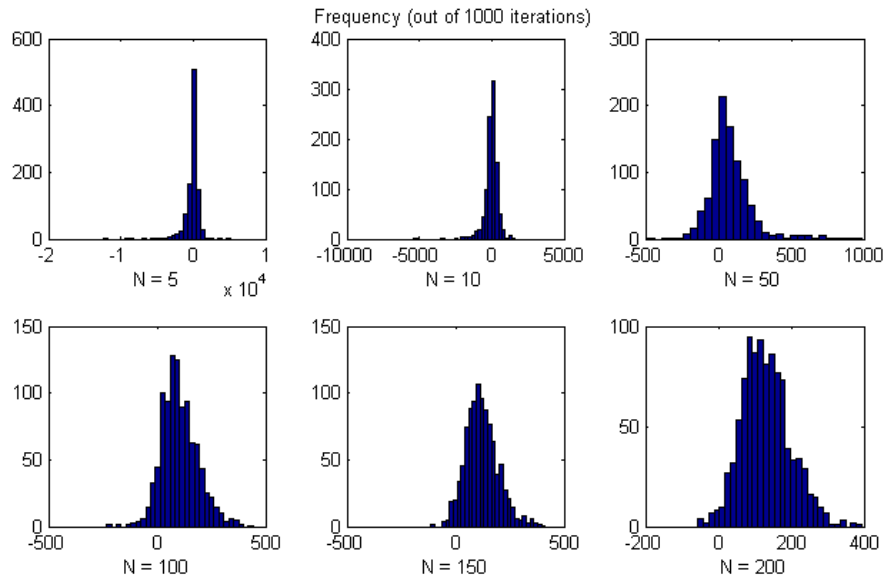


Figure 3 Average of the revenue difference: SAA - MaxMin Ratio

In Figure 3, we plot the distribution of the average difference in the SAA revenue minus the

MaxMin Ratio revenue. More specifically, for each sample size N , we sampled 1000 times the set of in-sample data of size N . Each one of these data sets is what we call an iteration of the experiment. For each of these iterations, we solved the pricing policies and evaluated the revenues for another 2000 out-of-sample data points and evaluated the difference in revenues between the policies for each of these points. This average difference obtained over those 2000 points correspond to one particular iteration. The histograms in Figure 3 show the distribution of this average difference obtained over all 1000 iterations of this experiment.

One can initially observe that for $N = 200$ the average difference lies clearly on the positive side of the histogram, showing that the SAA performs better on average than the MaxMin Ratio. When N is reduced, the histograms move closer to zero. When $N = 5$, the average difference seems to lie somewhere around zero, with a bias towards the negative side, implying that when the sample size is small the MaxMin Ratio will often perform better on average. These observations are confirmed by Table 1, where we display the mean average difference in the second column, with a 90% confidence interval (third column) calculated from the percentile of the empirical distribution formed over 1000 iterations. When $N = 5$ and $N = 10$, the average difference in revenues between the SAA and the Maxmin Ratio is in fact negative on the average of all iterations of the experiment, although not statistically significant, since the confidence interval contains zero. In other words, the confidence intervals suggest that there is not enough evidence to say the average difference in revenues is either positive or negative. On the other hand, when $N = 150$ or $N = 200$, we can see that with a 90% confidence level, the average difference is positive, meaning the SAA policy will provide better average revenues than the MaxMin Ratio.

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 other words, we have a statistically significant evidence that the SAA revenues will do much worse than the MaxMin Ratio in the bad cases, while not so much better in the better 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.

In the following experiments we will perform paired comparisons between the different robust models. In Figure 4, 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 4 Paired comparison between the MaxMin Ratio and the MaxMin

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

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 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




Figure 5 Paired comparison between the MaxMin Ratio and the MinMax Regret

to work.

In the next experiment, we compare the two regret-based robust models: MaxMin Ratio and MinMax Regret. Figure 5 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 suggests the MinMax Regret could be the more robust model between the two regret-based models.

3.2. 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. Consider the following parameters for the model: $T = 2$, $n = m = 1$, $a = [200, 200]$, $b = [0.5, 0.5]$, $C = 120$, $o = 1000$, $s = 10$. For simplicity, we will also consider a box-type uncertainty set, where $U = \{\delta : |\delta_t| \leq 15, \forall t = 1, 2\}$. In fact these are the same parameters as in the previous section. 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.

	Average		5% CVaR		95% CVaR	
N	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. 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, with the confidence intervals moving away from the positive side as the sample size increases. The CVaR comparisons, on the other hand, are inconclusive.

	Average		5% CVaR		95% CVaR	
N	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. Since the MaxMin is clearly dominated by the MaxMin Ratio, we will not expose here the comparison of the MaxMin with the other models.

	Average		5% CVaR		95% CVaR	
N	Mean	90% CI	Mean	90% CI	Mean	90% CI
5	-33	[-157,57]	-312	[-926,-50]	147	[15,479]
10	-35	[-133,51]	-434	[-1205,-155]	227	[47,582]
50	2	[-86,82]	-532	[-931,-376]	296	[124,419]
100	22	[-48,98]	-526	[-801,-415]	327	[232,417]
150	31	[-26,92]	-518	[-708,-446]	349	[266,409]
200	31	[-17,93]	-533	[-682,-434]	390	[308,410]

Table 6 Using artificial sampling distribution: Average and CVaR of the difference Revenue MaxMin Ratio - Revenue MinMax Regret

In the experiment of Table 6, we display the results for the revenues of the MaxMin Ratio model minus the MinMax Regret model. Note that the MinMax Regret model appears to obtain a better average difference than the MaxMin Ratio for sample sizes up to 10. As the sample size increases to 50 and above, the MaxMin Ratio performs better on average. Besides this initial evidence, either one of these results cannot be fully supported by the confidence intervals, which are containing zero in all instances of sample size. On the other hand, we notice that these confidence intervals are moving away from the negative side as the sample size increases, which agrees with our conclusion.


On the other hand, the CVaR comparison indicates that the MinMax Regret is in general more robust than the MaxMin Ratio, as the 5% better cases of the MinMax Regret (negative side) are on average larger than the 5% better cases of the MaxMin Ratio. This result is in fact supported by the confidence intervals, which do not overlap (in absolute value) for sample sizes of $N = 150$ and $N = 200$.

We should mention that comparing the MinMax Regret model with the SAA led to the same results as the comparison between MaxMin Ratio and the SAA and therefore these results are not displayed here as they do not add any additional insight.

To summarize the experiments performed using the artificial sampling distribution, where the samples provided come from a 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. When comparing the two regret type models, we have evidence that the MaxMin Ratio should perform better on average than MinMax Regret for large enough sample sizes, but the MinMax Regret is clearly more robust than the MaxMin Ratio.

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 departures 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 6 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 3.1 and 3.2. Instead, the concept of sample standard deviation will be used to express volatility of the revenue outcomes and therefore some form of robustness.

**Figure 6** Price x Demand Estimation

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 7 Airline case: Revenue performance for each pricing model

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 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 7.

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 network dynamic pricing problem using a sampling based approach to find approximate-closed-loop robust pricing policies.

Throughout the paper, we have shown how the adjustable robust pricing problem can easily turn into a non-convex problem, which is intractable to solve using traditional robust optimization techniques. Instead we proposed the sampling based approach where we can easily solve the sampled problem as a convex optimization model. One question that arises from this approach is how many samples do we need to have a performance guarantee on our scenario-based solution. To answer this question, 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. Some of the key managerial insights of our numerical experiments are:

- 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 uncertainty 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. This 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. As in the previous case, where the sampling distribution was the true distribution, the MinMax Regret appears to be more conservative than the MaxMin Ratio.

Note that the methodological framework presented in this paper is quite general and is not limited to dynamic pricing. For instance, Perakis (2009) developed a dynamic traffic assignment model using a very similar framework. Because of the practicality of these models, we believe these sampling based optimization models can easily be applied to many more application areas.

It would also be interesting to further research the Sample Average Approximation approach, by possibly finding a theoretical performance guarantee based on the number of samples used and comparing this with the bounds for the robust framework. We find that a specially promising direction of research will be to better understand the differences in behavior between the robust and the SAA models when using small sample sizes. Also, more research can be done on showing under which conditions can the affine policies achieve the optimal value of the fully closed-loop solution.

Acknowledgments

We are thankful for the multiple conversations and constructive comments of Gabriel Bitran, Marco Campi, Vivek Farias, Stephen Graves, David Morton, Joline Uichanco, Garrett van Ryzin and the participants of

the Fall 2008 Operations Management seminar at MIT. This research was partially supported by NSF grants 0556106-CMII and 0824674-CMII and the Singapore MIT Alliance Program.

References

- Adida, E., G. Perakis. 2005. A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematical Programming Special Issue on Robust Optimization* .
- Akan, M., B. Ata. 2009. Bid-price controls for network revenue management: Martingale characterization of optimal bid prices. *Mathematics of Operations Research* **34**(4) 912–936.
- Araman, V., R. Caldentey. 2009. Dynamic pricing for non-perishable products with demand learning. *Operations Research* (to appear).
- Bagnoli, M., T. Bergstrom. 1989. Log-concave probability and its applications .
- Ball, M. O., M. Queyranne. 2006. Toward robust revenue management: Competitive analysis of online booking. Robert H. Smith School Research Paper No. RHS 06-021. Available at SSRN: <http://ssrn.com/abstract=896547>.
- Bayraksan, G., D. P. Morton. 2006. Assessing solution quality in stochastic programs. *Mathematical Programming* **108**(2-3) 495–514.
- Belobaba, P. P. 1987. Air travel demand and airline seat inventory management. Ph.D. thesis, MIT.
- Belobaba, P. P. 1992. Optimal versus heuristic methods for nested seat allocation. *AGIFORS Reservations and Yield Management Study Group Meeting* .
- Ben-Tal, A., B. Golany, A. Nemirovski, J.P. Vial. 2005. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing and Service Operations Management* **7**(3) 248–271.
- Ben-Tal, A., A. Goryashko, E. Guslitzer, A. Nemirovski. 2004. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming* **99**(2) 351–376.
- Ben-Tal, A., A. Nemirovski. 1998. Robust convex optimization. *Mathematics of Operations Research* **23** 769–805.
- Ben-Tal, A., A. Nemirovski. 1999. Robust solutions to uncertain linear programs. *Operations Research Letters* **25** 1–13.
- Ben-Tal, A., A. Nemirovski. 2000. Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming* **88** 411–424.
- Bertsimas, D., S. de Boer. 2005. Simulation-based booking limits for airline revenue management. *Operations Research* **53**(1) 90–106.
- Bertsimas, D., D. Iancu, P. Parrilo. 2009. Optimality of affine policies in multi-stage robust optimization Working paper, MIT.
- Bertsimas, D., I. Popescu. 2003. Revenue management in a dynamic network environment. *Transportation Science* **37**(3).

- Bertsimas, D., M. Sim. 2004. The price of robustness. *Operations Research* **52**(1) 35–53.
- Bertsimas, D., A. Thiele. 2006. A robust optimization approach to supply chain management. *Operations Research* **54**(1).
- Besbes, O., A. Zeevi. 2008. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research* (to appear).
- Birbil, S. I., J. B. G. Frenk, J. A. S. Gromicho, S. Zhang. 2006. An integrated approach to single-leg airline revenue management: The role of robust optimization. *Econometric Institute Report E.I.* 2006-20.
- Bitran, G., R. Caldentey. 2003. An overview of pricing models and revenue management. *Manufacturing and Service Operations Management* **5**(3) 203–229.
- Bobkov, S. G. 1999. Isoperimetric and analytic inequalities for log-concave probability measures. *The Annals of Probability* (4).
- Calafiore, G. 2009. On the expected probability of constraint violation in sampled convex programs. *Journal of Optimization Theory and Applications* **143**(2) 405–412.
- Calafiore, G.C., M.C. Campi. 2005. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming* **102** 25–46.
- Calafiore, G.C., M.C. Campi. 2006. The scenario approach to robust control design. *IEEE Transactions On Automatic Control* **51**(5) 742–753.
- Campi, M.C., S. Garatti. 2008. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization* **19**(3) 1211–1230.
- Caramanis, C. 2006. Adaptable optimization: Theory and algorithms. PhD dissertation, Department of Electrical Engineering and Computer Science, MIT.
- Chiralaksanakul, A., D. P. Morton. 2004. Assessing policy quality in multi-stage stochastic programming. *Stochastic Programming E-Print Series (www.speps.org)*.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing: Research overview, current practices and future directions. *Management Science* **49** 1287–1309.
- Eren, S. S., C. Maglaras. 2009. Monopoly pricing with limited demand information. Working paper, Columbia University.
- Farias, V., B. Van Roy. 2009. Dynamic pricing with a prior on market response. *Operations Research* (to appear).
- Gallego, G., I. Moon. 1993. The distribution free newsboy problem: Review and extensions. *The Journal of the Operational Research Society* **44**(8).
- Goldfarb, D., G. Iyengar. 2003. Robust convex quadratically constrained programs. *Mathematical Programming* **97**(3) 495–515.

- Kleywegt, A., A. Shapiro, T. Homem-de Mello. 2001. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, **12**(2) 479–502.
- Lan, Y., H. Gao, M. Ball, I. Karaesmen. 2006. Revenue management with limited demand information. *Management Science* (to appear).
- Levi, R., G. Perakis, J. Uichanco. 2010. New bounds and insights on the data-driven newsvendor model Working paper, MIT.
- Levi, R., R. O. Roundy, D. B. Shmoys. 2007. Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research* **32**(4) 821–839.
- Lim, A.E.B., J.G. Shanthikumar. 2007. Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research* **55**(2) 198–214.
- Lim, A.E.B., J.G. Shanthikumar, T. Watwai. 2008. Robust multi-product pricing. Working paper, Department of Industrial Engineering and Operations Research, University of California, Berkeley, CA.
- Littlewood, K. 1972. Forecasting and control of passenger bookings. *Proceedings of the Twelfth Annual AGIFORS Symposium*.
- Nguyen, D. T., G. Perakis. 2005. Robust competitive pricing for multiple perishable products. Working paper, MIT.
- Pagnoncelli, B., S. Ahmed, A. Shapiro. 2009. Computational study of a chance constrained portfolio selection problem. *Journal of Optimization Theory and Applications* (to appear).
- Perakis, G. 2009. Alleviating travel delay uncertainties in traffic assignment and traffic equilibrium. *Proceedings for NSF Engineering Research and Innovation Conference*.
- Perakis, G., G. Roels. 2008. Regret in the newsvendor model with partial information. *Operations Research* **56**(1) 188–203.
- Perakis, G., A. Sood. 2006. Competitive multi-period pricing for perishable products. *Mathematical Programming*.
- Roels, G., G. Perakis. 2007. Robust controls for network revenue management. *Manufacturing & Service Operations Management* (to appear).
- Rusmevichientong, P., B. Van Roy, P. W. Glynn. 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., G. van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Science* **44**(11) 1577–1593.
- Talluri, K., G. van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers.
- Thiele, A. 2006. Single-product pricing via robust optimization. Working paper, Lehigh University.
- van Ryzin, G., J. McGill. 2000. Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science* **46**(6) 760–775. doi: <http://dx.doi.org/10.1287/mnsc.46.6.760.11936>.
- van Ryzin, G., G. Vulcano. 2008. Simulation-based optimization of virtual nesting controls of network revenue management. *Operations Research* **56**(4) 865–880.
- Yue, J., B. Chen, Wang M.-C. 2006. Expected value of distribution information for the newsvendor problem. *Operations Research* **54**(6) 1128–1136.
- Zhan, R., Z. Shen. 2005. Newsvendor problem with pricing: properties, algorithms, and simulation. *WSC '05: Proceedings of the 37th conference on Winter simulation*. Winter Simulation Conference, 1743–1748.
- Zhang, L. 2006. Multi-period pricing for perishable products: Uncertainty and competition. MS dissertation, Computation Design and Optimization, MIT.

Appendix A: Notation.

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

- C : Vector of initial inventories for each resource.
- T : Length of the time horizon.
- δ : Vector of the realization of the demand uncertainty, containing one component for each product and time.
- U : Uncertainty set, $\delta \in U$.
- s : Vector of pricing decisions.
- S : Strategy space, independent of the demand uncertainty.
- $p_{j,t}(s, \delta)$: Price of product j at time t as a function of the pricing strategy and the uncertainty realized up to time t .
- $\tilde{d}_{j,t}(s, \delta)$: Realized demand of product j at time t , given the pricing strategy and the uncertainty.
- o : Vector of overbooking penalties per unit of resource.
- g : Vector of salvage values per unit of resource.
- $\tilde{w}(s, \delta)$: Vector function that maps the pricing strategy and the realized uncertainty into overbooked/left-over inventory.
- w : Vector of slack variables to account for overbooked/left-over inventory. At optimality, $w_l = \tilde{w}_l(y, \delta)$, while w^+ is the amount of overbooked units and w^- is the left-over inventory.
- $\Pi(s, \delta)$: Actual revenue given a pricing policy and realized uncertainty.
- $\Pi^*(\delta)$: Perfect hindsight revenue, given some uncertainty vector.
- $h^{obj}(s, \delta)$: 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.
- $f(x, \delta)$: Short notation for all 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 .
- k : Bounded likelihood ratio factor between the sampling distribution \mathcal{P} and the true distribution \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 function $h(u, v, \delta_1, \delta_2)$ for any fixed u and v is:

$$h(\delta_1, \delta_2) = u_1(a_1 - b_1 u_1 + \delta_1) + (u_2 + v_{2,1} \delta_1)(a_2 - b_2(u_2 + v_{2,1} \delta_1) + \delta_2)$$

The eigenvalues of the Hessian matrix of this function are $(-b_2 v_{2,1} + \sqrt{b_2^2 v_{2,1}^2 + 2})v_{2,1}, (-b_2 v_{2,1} - \sqrt{b_2^2 v_{2,1}^2 + 2})v_{2,1}$. These will always have one positive and one negative eigenvalue. Therefore the function will neither be convex nor concave.

Appendix C: Existence of an optimal solution.

In this section we prove there exists an optimal solution to (5). Note that the function $f(s, \delta)$ 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 $\bar{f}(s) = \min_{\delta \in U} f(s, \delta)$ is also a continuous function of s . For some $\Delta > 0$:

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

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

$$f(s + \Delta, \delta_1) - f(s, \delta_1) \leq \bar{f}(s + \Delta) - \bar{f}(s) \leq f(s + \Delta, \delta_2) - f(s, \delta_2)$$

By letting $\Delta \rightarrow 0$, both sides will go to 0. It follows that:

$$\lim_{\Delta \rightarrow 0} |\bar{f}(s + \Delta) - \bar{f}(s)| = 0$$

Which implies that $\bar{f}(s)$ is a continuous function. Since S is a compact set, we use Weierstrass Theorem to guarantee existence of solution to (5).

Appendix D: Proof of Theorem 1.

Before proving Theorem 1, we first introduce Lemma D.1 below that bounds the density of the mode of a log-concave distribution.

Lemma D.1 *Let f be the density of a one-dimensional random variable X with standard deviation σ and mode M .*

- (a) *If f is log-concave, then $f(M) \leq \frac{e^{\sqrt{6}}}{\sigma\sqrt{2}}$.*
- (b) *If f is log-concave and symmetric, then $f(M) \leq \frac{1}{\sigma\sqrt{2}}$.*

Let F be the cumulative density of the random variable X and m be the median of X . Define the isoperimetric constant as $Is = \inf_x \frac{f(x)}{\min\{F(x), (1-F(x))\}}$, see Bobkov (1999) for more details about the properties of the isoperimetric constant.

For any $x \leq m$, we have that $F(x) \leq 1 - F(x)$. Therefore, $Is \leq \frac{f(x)}{F(x)}$ and from the log-concavity of f , we know that the reversed hazard rate $\frac{f(x)}{F(x)}$ is decreasing, then the minimum will be achieved at the median. Similarly for the other side, for any $x \geq m$, we have that $F(x) \geq 1 - F(x)$. Therefore, $Is \leq \frac{f(x)}{1-F(x)}$ and from the log-concavity of f , we have an increasing failure rate $\frac{f(x)}{1-F(x)}$. Again, the minimum will be achieved at the median. Since at the median $F(m) = 1 - F(m) = 1/2$, then $Is = \frac{f(m)}{F(m)} = 2f(m)$.

From Bobkov (1999), Proposition 4.1, we have that $Is^2 \leq \frac{2}{\sigma^2}$. For symmetric distributions the median is equal to the mode, which directly implies the result for the symmetric case of Lemma D.1(b):

$$(2f(M))^2 = (2f(m))^2 \leq \frac{2}{\sigma^2} \Rightarrow f(M) \leq \frac{1}{\sigma\sqrt{2}}$$

For non-symmetric distributions, Lemma D.1(a), we must find a relation between $f(M)$ and $f(m)$. Note that since $\log f(\cdot)$ is a concave function, the first order approximation will always upper-bound the original function (as seen in Lemma 3.3 of Levi et al. (2010)). Therefore:

$$\log f(M) \leq \log f(m) + \frac{f'(m)}{f(m)}(M - m)$$

For non-differentiable distributions, note this relation is true for all sub-gradients $f'(m)$ at m . From the decreasing reversed hazard rate property of log-concave distributions, we have that: $\frac{d}{dx} \frac{f(x)}{F(x)} \leq 0$, which implies that $\frac{f'(x)}{f(x)} \leq \frac{f(x)}{F(x)}$. When $M \geq m$, we get:

$$\log f(M) \leq \log f(m) + \frac{f(m)}{F(m)}(M - m)$$

From the increasing failure rate property of log-concave distributions, we have that $\frac{d}{dx} \frac{f(x)}{1-F(x)} \geq 0$, which implies $\frac{f'(x)}{f(x)} \geq -\frac{f(x)}{1-F(x)}$. Then for the case $M < m$, we have:

$$-\log f(M) \geq -\log f(m) + \frac{f'(m)}{f(m)}(m - M) \geq -\log f(m) - \frac{f(m)}{1-F(m)}(m - M)$$

Combining the relations above and the fact that $F(m) = 1 - F(m) = 1/2$, we obtain for either $M \geq m$ or $M < m$:

$$\log f(M) \leq \log f(m) + 2f(m)|M - m|$$

Since any log-concave distribution is unimodal, we can use the relations in Basu and DasGupta (1992), Corollary 4, to bound the distance between the mode and the median: $|M - m| \leq \sigma\sqrt{3}$. Therefore, $\log f(M) \leq \log f(m) + 2f(m)\sigma\sqrt{3}$. Since $(2f(m))^2 \leq \frac{2}{\sigma^2}$, we have that $f(m) \leq \frac{1}{\sigma\sqrt{2}}$, which implies $\log f(M) \leq \log \frac{1}{\sigma\sqrt{2}} + 2\frac{1}{\sigma\sqrt{2}}\sigma\sqrt{3} = \log \frac{1}{\sigma\sqrt{2}} + \sqrt{6}$. Taking the exponential on both sides we get $f(M) \leq \frac{e\sqrt{6}}{\sigma\sqrt{2}}$, which concludes the proof of Lemma D.1(a). \square

Using Lemma D.1, we can directly obtain the likelihood ratio bound k on Theorem 1 by using a bound on the density of the mode $f(M)$ for log-concave distributions. Note that the uniform sampling distribution has a constant density value of $1/(ub - lb)$ for each product $j = 1, \dots, n$ and time period $t = 1, \dots, T$, which implies a ratio bound of $\left[(ub - lb) \frac{e\sqrt{6}}{\sigma\sqrt{2}} \right]$ for each product and time. Using the independence property, we can just multiply this bound nT times to get the desired result. The symmetric case follows from the same argument.

Appendix E: 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 $V_{\mathcal{Q}}(x^N) \doteq \mathbf{P}_{\mathcal{Q}}\{\delta \in \Delta : f(x^N, \epsilon) > 0\}$.

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 $B_I \doteq \{(\delta^{(1)}, \dots, \delta^{(N)}) :$

$V_Q(x^I) > \epsilon\}$. 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 $V_Q(x^N) > \epsilon$ 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 $V_Q(x^I) > \epsilon$ 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 $f(x^I, \delta) \leq 0$, 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:

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

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

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

□

Appendix F: Sample size bound for sampling with true distribution

The following theorem states the confidence level bound for the data-driven case where the sampling distribution is the true distribution.

Theorem F.1 *From Campi and Garatti (2008), if we are given the scenarios of uncertainty $\delta^{(1)}, \dots, \delta^{(N)}$ drawn from the true uncertainty distribution \mathcal{Q} . Let n_x be the dimension of the strategy space X and ϵ be the robust level parameter. Then define a “risk of failure” parameter $\beta(N, \epsilon)$ as:*

$$\beta(N, \epsilon) \doteq \sum_{i=0}^{n_x-1} \binom{N}{i} \epsilon^i (1 - \epsilon)^{N-i} \quad (\text{F.1})$$

Then with probability greater than $(1 - \beta(N, \epsilon))$, the solution found using this method is ϵ -level robustly feasible, i.e.,

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

If we are given a desired confidence level β and a robust level ϵ , it is natural to ask what is the minimum number of samples required. This design question can be answered by numerically solving for N the equation in (F.1). Calafiore (2009) has shown that by using the Chernoff bound on the lower binomial tail, we can obtain an explicit formula for N that will guarantee that (F.1) is satisfied.

Corollary F.1 *From Calafiore (2009), if the sample size N follows:*

$$N \geq \frac{2}{\epsilon} \left(\ln \frac{1}{\beta} + n_x \right) + 1$$

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

Numerically testing the bound:

The goal of this experiment is to validate the results of Theorem F.1, which is provably a tight confidence/robustness bound for the data-driven robust problem. The bound from Theorem 2 is less tight, i.e., requires more samples to achieve the same level of confidence. For this reason, it is sufficient for us to test the tighter bound of Theorem F.1 in order to verify the effectiveness of both these bounds. Note that this experiment is done under the assumption that the data samples provided are from the true distribution of uncertainty, the data-driven case.

To test the bound we consider a small instance of the MaxMin robust model, with a single product ($n = m = 1$), two time periods ($T = 2$), a linear demand model and affinely adjustable pricing policies. Using an adjustable pricing policy, the set of pricing decision variables has dimension $n_x = 4$, with one static price for each time period u_1, u_2 , an adjustable component v_2 and the robust level z . Note that we will use dummy variables for the net inventory $w^{(i)}$ that are completely determined by the other variables, therefore they do not count for the dimension of the strategy space in our confidence/robust bound calculation. More specifically, the demand at each period is given as a linear function of the price, where a is a constant demand term, b is the price elasticity of demand and δ is an additive uncertain component: $\tilde{d}_t(u, v, \delta) = a_t - b_t(u_t + v_t \sum_{k=1}^{t-1} \delta_k^{(i)}) + \delta_t$. For simplicity of notation, we will not separate the case for the first time period, $t = 1$, where there is no adjustability to be considered. Therefore, wherever you consider the case of $t = 1$, define the summation from $k = 1$ to $t - 1$ as an empty sum, with zero value. Moreover, we will assume δ lies within a box-type uncertainty set, which is widely used in the robust optimization literature, defined as: $U = \{(\delta_1, \delta_2, \dots, \delta_T) : |\delta_t| \leq \Gamma_t, \forall t = 1, \dots, T\}$. The experiments in this section can be performed with any type of convex uncertainty set.

The parameters used for this test were $T = 2$, $a = [200, 200]$, $b = [0.5, 0.5]$, $C = 100$ and $\Gamma = [15, 30]$. For simplicity, we will consider a model virtually without overbooking or salvage value: $o = 10000, s = 10$. We assume for this test that the uncertainty is uniformly distributed over the uncertainty set. Moreover, we are sampling our data points from the same distribution. In other words, we are testing the data-driven case, where the scenario samples $\delta^{(1)}, \dots, \delta^{(N)}$ come from the true distribution, possibly from historical data.

At each experiment, we draw $N = 200$ uncertainty scenarios to use in the optimization model to determine the pricing policy and another 2000 scenarios to evaluate the out-of-sample performance of this sampling-based pricing policy. Define the estimated probability of violation \hat{V} as the proportion of the out-of-sample scenarios for which the constraints are violated under the scenario based policy.

We repeated the experiment 300 times in order to obtain multiple instances of the problem, each with a different probability of violation. In other words, we drew other samples of $N = 200$ scenarios and repeated




Figure 7 Bound Test for $N = 200$: $\epsilon \times \beta$

the procedure described before. We counted the probability of violation \hat{V} for each experiment and ordered the experiments in a decreasing vector. For any given failure level $\hat{\beta}$, define the $\hat{\beta}$ -quantile of this vector as the estimated robust level $\hat{\epsilon}$ associated with $\hat{\beta}$. This means that a proportion $\hat{\beta}$ of the experiments we ran, have demonstrated a probability of violation greater than $\hat{\epsilon}$.

$$\mathbf{P}(\hat{V} > \hat{\epsilon}) \doteq \hat{\beta}(\hat{\epsilon})$$

From Theorem F.1, we also have a theoretical guarantee on this “risk of failure” β :

$$\mathbf{P}(V_Q > \epsilon) \leq \beta(N, \epsilon) \doteq \sum_{i=0}^{n_x-1} \binom{N}{i} \epsilon^i (1 - \epsilon)^{N-i}$$

Using $N = 200$ samples, Figure 7 shows, with a dashed line, the theoretical bound on the ϵ -robust level (y-axis) for each possible failure level β (x-axis). The solid line shows the estimated $\hat{\epsilon}$ for the same failure levels β . For example, one should read this graph as follows: for $\beta = 0.0395$, i.e., confidence level of 96.05%, the solution for the MinMax model is 4%-robustly feasible, according to Theorem F.1. In other words, assuming we have a good sample, there is at most a 4% chance that the actual competitive ratio will be lower than the solution provided by scenario based solution. When observing the number of violations that actually happened, the estimated probability of violation, 3.96% of the simulations had a ratio below the sample based ratio. At confidence level 99.97%, the theoretical robust level is 7% and our estimated probability of violation was 6.05%. Note how the estimated probabilities of violation are fairly close to the theoretical bound. In Figure 7, we present the outcome of the test with 300 experiments, but it is interesting to note that the empirical line converges to the theoretical bound roughly within 100 experiments. Overall, we have a strong empirical evidence that the bound in Theorem F.1 is effective, if not in fact tight, for the pricing problem.

Appendix G: Summary of optimization models in the numerical experiments.

Throughout the modeling section of the paper, Section 2, we described the general framework of the pricing models using simplified concise notation. In this section of the appendix, we will summarize the full model

instances used to run the simulations in Section 3, trying to be as clear as possible in case the reader wishes to implement any of these models in a convex optimization solver. Define the following model as the **MaxMin**:

$$\begin{aligned} & \max_{z, u, v, p, w, w^+} z \\ & s.t. \left\{ \begin{array}{l} z \leq \sum_{t=1}^T p_t^{(i)} (a_t - b_t p_t^{(i)} + \delta_t^{(i)}) - \sum_{l=1}^m ((o_l - g_l) w_l^{+(i)} + g_l w_l^{(i)}) \\ \sum_{t=1}^T M(a_t - b_t p_t^{(i)} + \delta_t^{(i)}) \leq C + w^{(i)} \\ p_t^{(i)} = u_{j,t} + \sum_{l=1}^m v_{j,l,t} \sum_{j'=1}^n \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}^{(i)} \geq 0, \\ w_l^{(i)} \leq w_l^{+(i)}, \\ w_l^{+(i)} \geq 0, \end{array} \right. \quad \begin{array}{l} \forall t = 1, \dots, T \\ \forall l = 1, \dots, m \\ \forall l = 1, \dots, m \end{array} \end{array} \quad \forall i = 1, \dots, N \quad (\text{G.1})$$

For the regret based models, define the perfect hindsight revenue as the solution of the following model:

$$\begin{aligned} \Pi^*(\delta) = & \max_{u, v, p, w, w^+} \sum_{t=1}^T p_t (a_t - b_t p_t + \delta_t) - \sum_{l=1}^m ((o_l - g_l) w_l^+ + g_l w_l) \\ & s.t. \sum_{t=1}^T M(a_t - b_t p_t + \delta_t) \leq C + w \\ & p_t \geq 0, \quad \forall t = 1, \dots, T \\ & w_l \leq w_l^+, \quad \forall l = 1, \dots, m \\ & w_l^+ \geq 0, \quad \forall l = 1, \dots, m \end{aligned} \quad (\text{G.2})$$

Using the solution of the hindsight revenue $\Pi^*(\delta^{(i)})$ for every sample of $\delta^{(i)}$ provided, we can define the **MixMax Regret** model as:

$$\begin{aligned} & \max_{z, u, v, p, w, w^+} z \\ & s.t. \left\{ \begin{array}{l} z \leq \sum_{t=1}^T p_t^{(i)} (a_t - b_t p_t^{(i)} + \delta_t^{(i)}) - \sum_{l=1}^m ((o_l - g_l) w_l^{+(i)} + g_l w_l^{(i)}) - \Pi^*(\delta^{(i)}) \\ \sum_{t=1}^T M(a_t - b_t p_t^{(i)} + \delta_t^{(i)}) \leq C + w^{(i)} \\ p_t^{(i)} = u_{j,t} + \sum_{l=1}^m v_{j,l,t} \sum_{j'=1}^n \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}^{(i)} \geq 0, \\ w_l^{(i)} \leq w_l^{+(i)}, \\ w_l^{+(i)} \geq 0, \end{array} \right. \quad \begin{array}{l} \forall t = 1, \dots, T \\ \forall l = 1, \dots, m \\ \forall l = 1, \dots, m \end{array} \end{array} \quad \forall i = 1, \dots, N \quad (\text{G.3})$$

Similarly, the **MaxMin Ratio** model is defined as:

$$\begin{aligned} & \max_{z, u, v, p, w, w^+} z \\ & s.t. \left\{ \begin{array}{l} z \leq \left(\sum_{t=1}^T p_t^{(i)} (a_t - b_t p_t^{(i)} + \delta_t^{(i)}) - \sum_{l=1}^m ((o_l - g_l) w_l^{+(i)} + g_l w_l^{(i)}) \right) / \Pi^*(\delta^{(i)}) \\ \sum_{t=1}^T M(a_t - b_t p_t^{(i)} + \delta_t^{(i)}) \leq C + w^{(i)} \\ p_t^{(i)} = u_{j,t} + \sum_{l=1}^m v_{j,l,t} \sum_{j'=1}^n \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}^{(i)} \geq 0, \\ w_l^{(i)} \leq w_l^{+(i)}, \\ w_l^{+(i)} \geq 0, \end{array} \right. \quad \begin{array}{l} \forall t = 1, \dots, T \\ \forall l = 1, \dots, m \\ \forall l = 1, \dots, m \end{array} \end{array} \quad \forall i = 1, \dots, N \quad (\text{G.4})$$

Alternatively, we will also use the **Sample Average Approximation (SAA)** model as a benchmark, which can be defined as:

$$\begin{aligned}
 & \max_{z, u, v, p, w, w^+} \frac{1}{N} \sum_{i=1}^N \left[\sum_{t=1}^T p_t^{(i)} \left(a_t - b_t p_t^{(i)} + \delta_t^{(i)} \right) - \sum_{l=1}^m \left((o_l - g_l) w_l^{+(i)} + g_l w_l^{(i)} \right) \right] \\
 & s.t. \quad \left\{ \begin{array}{l} \sum_{t=1}^T M \left(a_t - b_t p_t^{(i)} + \delta_t^{(i)} \right) \leq C + w^{(i)} \\ p_t^{(i)} = u_{j,t} + \sum_{l=1}^m v_{j,l,t} \sum_{j'=1}^n \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}^{(i)} \geq 0, \forall t = 1, \dots, T \\ w_l^{(i)} \leq w_l^{+(i)}, \quad \forall l = 1, \dots, m \\ w_l^{+(i)} \geq 0, \quad \forall l = 1, \dots, m \end{array} \right\} \quad \forall i = 1, \dots, N \quad (G.5)
 \end{aligned}$$

Appendix H: Static vs. Adjustable policies.

In Figure 8, we display the out-of-sample revenues between the static and the adjustable policies for the MinMax robust model. The static counterpart is defined the same way as in (11), but with the price function consisting only of the static component $u_{j,t}$. We can clearly see that the static policy is dominated by the adjustable policy, which performs a better worst case revenue as well as in overall revenue distribution. This observation is exactly as we expected, since the adjustable policy gives the model an extra degree of freedom, which should only improve the performance of the model.




Figure 8 Example of out-of-sample revenue distribution: Static vs. Adjustable policies

Appendix I: Network example: multiple products and resources.

For this experiment, consider a simple network problem, with $m = 2$ resources and $n = 3$ products. The incidence matrix used is:

$$M = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}$$

This can be viewed as the airline problem with two legs who is selling either the single leg itineraries or a combination of both legs. We will again use linear demand functions with additive uncertainty for simplicity. The reader should note that this is used here for simplicity of notation only and the use of other demand and

uncertainty models can be easily incorporated. We further assume the pricing policies for each product/time-period to be affine functions of the accumulated uncertainty for the corresponding resources consumed over the previous time-periods. For a given uncertainty realization $\delta^{(i)}$ and a set of pricing decisions u and v , define the prices at each period according to (3). Instead of using more complex policies, which could include cross-product and time dependencies, we use here a simpler policy as suggested in (3) (i.e. a policy for each product that depends only on its own remaining capacity) because we simply want to illustrate how to apply this framework in a network problem.

Note that the $\sum_{j'=1}^n \sum_{t'=1}^{t-1} M_{l,j'} \delta_{j',t'}^{(i)}$ is the amount of unpredicted resources of type l consumed by time t due to the realized uncertainty $\delta^{(i)}$ up to time $t-1$. The adjustable component of the pricing strategy, $v_{j,l,t}$ is the price impact on product j at time t from the consumption of resource l .

We compare the trade-offs between the worst-case performance and the average performance for each of these four models. The particular instance of the problem used in this experiment is: $T = 2$, $a_{j,t} = 200$, $b_{j,t} = 0.5$ for all products j and time t , $C_l = 120$, $o_l = 1000$, $g_l = 10$ for each resource l . These numbers were chosen to be somewhat similar to the case study displayed in Section 4. We defined a box type uncertainty set, $U = \{(\delta_1, \delta_2) : |\delta_t| \leq 15, \forall t = 1, 2\}$, and sampled data points using a normal distribution with zero mean and standard deviation of 15 and then truncated over the uncertainty set U . For the optimization part, we used 200 samples to obtain pricing policies for each method and simulated the out-of-sample performance of these policies using 2000 samples. Figure 9 displays the revenue distribution for an instance of this problem with our four models. The difference between the robust models does not seem significant, while the SAA presents a better average and lower worst case revenues. In general, the results for the network cases are consistent with the single product cases displayed in Section 3.




Figure 9 Example of out-of-sample revenue distribution for the network problem.




Figure 10 Logarithmic Demand Model: $T = 2, C = 10, a_t = 10, b_t = 2$

Appendix J: Nonlinear demand models.

Note that the numerical experiments of Section 3 were done using only linear demand functions, although our framework can also deal with nonlinear demands. In this section of the appendix, we simply demonstrate an application of the sampling based framework with a logarithmic demand function defined as $\tilde{d}(p) = a - b \log(p) + \delta$, where the demand is a linear function of the logarithm of the price plus some deviation δ . Using this demand function in a single product problem, we can apply the same underlying models defined in the previous section. For example, define the **MinMax Regret** as:

$$\begin{aligned}
 & \max_{p, u, v, z} z \\
 s.t. \quad & z \leq \sum_{t=1}^T p_t^{(i)} (a_t - b_t \log(p_t^{(i)}) + \delta_t^{(i)}) - o(w^{(i)})^+ - g(w^{(i)})^- - \Pi^*(\delta^{(i)}) & \forall i = 1, \dots, N \\
 & \sum_{t=1}^T (a_t - b_t \log(p_t^{(i)}) + \delta_t^{(i)}) \leq C + w^{(i)} & \forall i = 1, \dots, N \\
 & p_t^{(i)} = u_t + v_t \sum_{k=1}^{t-1} \delta_k^{(i)} \geq 0, & \forall t = 1, \dots, T, \forall i = 1, \dots, N
 \end{aligned}$$

Figure 10 displays the overall performance of the four different models under logarithmic demand. As we have shown in this section, there is no particular difference in the methodology needed to implement different demand models, as long as they satisfy Assumption 1. If we are provided with historical data of price and demand realizations, the demand function should be chosen according to the best possible fit to the data set and the uncertainty realizations will be the estimation errors.