Robust Timing of Markdowns

Michael Dziecichowicz* Daniela Caro† Aurélie Thiele‡

Revised: December 2014

Abstract

We propose an approach to the timing of markdowns over a finite time horizon in a continuous setting that does not require the precise knowledge of the underlying probabilities, instead relying on *range forecasts* for the arrival rates of the demand processes, and that captures the degree of the manager's risk aversion through intuitive *budget of uncertainty* functions. These budget functions bound the cumulative deviation of the arrival rates from their nominal values over the lengths of time for which a product is offered at a given price. A key issue is that using lengths of time as decision variables introduces non-convexities when budget functions are concave. In the single-product case, we describe a tractable and intuitive framework to incorporate uncertainty on customers' arrival rates, formulate the resulting robust optimization model, describe an efficient procedure to compute the optimal sale times, and provide theoretical insights. We then describe how to use the solution of the static robust optimization model to implement a dynamic markdown policy. We also extend the robust optimization approach to multiple products and suggest the idea of constraint aggregation to preserve performance for this type of problem structure.

1 Introduction

We analyze the problem of determining optimal sale times for products subject to demand uncertainty over a finite time horizon in a continuous setting. Taking an example from Gallego and van Ryzin (1994), consider a fashion retailer who is bringing to market a new line of clothing. The entire production process takes from six to eight months to complete, yet the firm plans to sell the entire inventory

*Manager, Ernst & Young, New York, NY

†Pricing Analyst, MillerCoors, Chicago, IL

‡Associate Professor, Lehigh University, Bethlehem, PA, aut204@lehigh.edu Corresponding author.

in as little as nine weeks. Because this is a new line, little is known about customer response to the price. The company has no resupply option during the sales season and products left over at the end of the time horizon have no resale value. Demand is uncertain but is influenced by price, and the merchandise manager must adjust the price throughout the selling season to maximize revenue. Gallego and van Ryzin (1994) also show that pricing policies that aim to run out of stock do not necessarily maximize revenue. Also, they prove that a unique price change is an asymptotically optimal policy as the volume of sales increases using dynamic programming.

Feng and Gallego (1995) show that it is optimal to set and adjust prices as soon as the time-to-go falls below a threshold that depends on the number of items currently in stock. That work assumes a Poisson process for the demands at each price point, and prices are defined to fall within a finite set for easier implementation. Also using dynamic programming, they show the importance of the timing of the markdown, and argue that an early near-optimal markdown typically results in higher revenues than a late, but optimal markdown. In some situations, dynamic programming might lead to policies that frequently change prices, which makes it difficult to segment the market and could confuse potential customers regarding the quality of the product. Also, retailers often wish to put multiple items on sale at the same time. Despite different demand forecasts for different products at different store locations, the merchandise manager must decide on the optimal time to begin the sale and on a single optimal markdown level to offer each product at across all locations. Because of its use of dynamic programming, the model in Feng and Gallego (1995) faces tractability issues when extended to the multi-product, multi-location case.

The reader is referred to Phillips (2005) and Elmaghraby and Keskinocak (2003) for an overview of markdown management and of the basics of pricing optimization, as well as a classification of the vast amount of existing literature dating back to Whitin (1955). Further, the reader is referred to Bitran and Caldentey (2003) for a survey on dynamic pricing models in revenue management. Bertsimas and Perakis (2006) present an optimization approach to dynamically set prices in order to maximize revenue in both competitive and non-competitive settings, and suggest that a decision-maker does better by incorporating realized demand information as time evolves into the policy. Adida and Perakis (2006) study a robust optimization approach to dynamic pricing in a multiple product setting under demand uncertainty, highlight the difficulties of computing the optimal pricing policy over a given time horizon, and suggest methods of keeping the formulation tractable. Robust optimization is a framework to handle uncertainty where the decision-maker optimizes the worst-case objective, with the worst case measured over a set centered at the nominal value of the uncertain parameters. In line with Bertsimas and

Sim (2004), the size of the uncertainty set is determined by a non-negative parameter called the *budget of uncertainty*; the higher the budget, the bigger the set and the worse the objective. The robust optimization literature up to 2008 is reviewed in Bertsimas et al. (2011). Gabrel et al. (2014) provides an overview of advances in robust optimization after 2008.

In this paper, we argue that robust optimization is well-suited for the problem of optimizing sale times under demand uncertainty because it offers intuitive modeling techniques as well as tractable formulations that can be solved efficiently. Our approach differs from the other frameworks available in the literature because we focus on applying robust optimization to *timing* decisions rather than prices. In our framework, the manager selects prices from an available menu (for instance offering a 10%, 25% and 40% discount); most importantly, he decides *when* to implement the price changes, in a continuous-time setting. Since the total demand at a given price depends on how long the item is offered at that price, our modeling framework requires us to apply robust optimization techniques at the level of *demand rates* rather than the demand itself, which creates new challenges in developing tractable robust formulations. (The fact that the budgets of uncertainty are now functions of the decision variables – the sale times – creates non-linearities. This will be elaborated upon below.)

A possible disadvantage of traditional robust optimization is that the optimal policy is static, in the sense that it does not incorporate new information over time, so that the model must be resolved to capture changes in inventory. As argued in Bitran and Caldentey (2003), the rapid evolution of information technologies and the corresponding growth of the Internet and e-commerce makes a static assumption potentially costly to decision-makers. In today's markets, it is possible to collect valuable information about demand, inventory levels, competitors' strategies, and to process it in real time. In such settings, decision-makers should react dynamically to changes in the marketplace. We explain how to use robust optimization to determine the optimal parameters of such policies in the setting at hand.

Contributions.

- We model uncertainty on the arrival rates of the demand processes through range forecasts and capture the manager's risk aversion through a "budget of uncertainty" function, which limits the cumulative deviation of the arrival rates from their mean and is determined by the decision-maker.

- In the nominal case and in the case where the budget of uncertainty function is linear in time, we provide closed-form solutions for the optimal sale time in the single-product case.

- In the case where the budget of uncertainty is concave and increasing, again for a single product, we derive a mixed-integer problem (MIP) that approximates the robust non-convex formulation.
- We develop a policy about the optimal time to put products on sale, which depends on both the number of items unsold and on the time-to-go, and use our robust optimization model to determine its parameters.
- We extend our analysis to the case of multiple products. In particular, we present the idea of constraint aggregation to maintain the performance of robust optimization for that problem structure.
- We provide numerical experiments to test the performance of the robust optimization approaches described in this paper.

Outline. We study the robust markdown problem in the case of a single product in Section 2 and provide the extension to multiple products in Section 3. Section 4 contains concluding remarks.

2 The Single-Product Case

We use the following notation, in line with Feng and Gallego (1995):

T:	length of the time horizon,
I:	number of possible markdowns,
s_i:	time of the i^{th} markdown (in $[0, T]$),
K:	number of items initially in inventory,
p_i:	price over time interval $[s_{i-1}, s_i)$,
λ_i:	arrival rate of demand at price p_i,
$N_i(s_i - s_{i-1})$:	demand at price p_i on $[s_{i-1}, s_i)$.

Throughout the paper, we assume that $p_1 > p_2 > \cdots > p_I$, by definition of markdowns. (As prices decrease, arrival rates increase.) Note that in our approach, the p_i are set parameters; the decision variables are the sales times s_i, where the decision maker switches from price p_i to p_{i+1}. We also assume that demand at a given price depends on the length of time for which the item is offered at that price, but not on the start and end times themselves. This is in line with the assumption of Poisson process in the stochastic optimization literature.

2.1 One sale time

2.1.1 Problem setup

Consider the case where the decision-maker has one possible sale time, $s \in [0, T]$. Let p_1, resp. p_2, be the price before, resp. after the item is put on sale ($p_1 > p_2$). In the stochastic version of this problem, the arrival rates λ_1 and λ_2 are assumed constant, and the random arrival process at price p_1, resp. p_2, satisfies: $E[N_1(s)] = \lambda_1 \, s$, resp. $E[N_2(T-s)] = \lambda_2 \, (T-s)$.

The robust optimization approach incorporates uncertainty on the demand processes by modeling the arrival rates as uncertain. Specifically, the arrival rate at time τ of the first process, denoted $\lambda_1(\tau)$, belongs to the range forecast $[\overline{\lambda}_1 - \widehat{\lambda}_1, \, \overline{\lambda}_1 + \widehat{\lambda}_1]$, where $\overline{\lambda}_1$ is the nominal value of the arrival rate and $\widehat{\lambda}_1$ is the half-width of its confidence interval, with $\widehat{\lambda}_1 < \overline{\lambda}_1$. This can be rewritten using the scaled deviations $z_1(\tau)$:

$$\lambda_1(\tau) = \overline{\lambda}_1 + \widehat{\lambda}_1 \, z_1(\tau), \; \forall \tau \leq s, \tag{1}$$

with $|z_1(\tau)| \leq 1$ for all $\tau \leq s$. We assume that the instantaneous scaled deviations $z_1(\tau)$ are independent of each other, which corresponds to a situation where customers do not exchange information over a product once the season has started. To limit the degree of conservatism of the solution, we bound the cumulative scaled deviation of the arrival process from its mean:

$$\int_0^s |z_1(\tau)| d\tau \leq \Gamma(s),$$

where $\Gamma : [0, T] \to \mathcal{R}^+$ is a concave, increasing function called the *budget of uncertainty* function, such that $\Gamma(0) = 0$, $\Gamma(s) \leq s$ for all $0 \leq s \leq T$ and $s \to s - \Gamma(s)$ is increasing. The fact that the Γ function is increasing with time reflects that longer time periods create more uncertainty, and the fact that it is concave reflects that independent sources of uncertainty tend to cancel each other out over time, in the spirit of the law of large numbers. The fact that $s \to s - \Gamma(s)$ is increasing enforces the intuitive property that the worst-case number of arrivals increases with time in the extreme worst case where $\widehat{\lambda}_i = \overline{\lambda}_i$, because the worst-case number of arrivals is then equal to $\overline{\lambda}_i(s - \Gamma(s))$, $i = 1, 2$. (This combined with Γ increasing in turn ensures that the property holds for less stringent values of $\widehat{\lambda}_i$.)

Furthermore, the uncertain cumulative demand at full price p_1, given that the sale begins at time s, is given by:

$$N_1(s) = \int_0^s \lambda_1(\tau) d\tau.$$

Reinjecting Eq. (1) yields:

$$N_1(s) = \overline{\lambda}_1\, s + \widehat{\lambda}_1 \int_0^s z_1(\tau) d\tau, \tag{2}$$

with $z_1 \in \mathcal{Z}_1$ where $\mathcal{Z}_1$ is defined by:

$$\mathcal{Z}_1 = \left\{ z_1 \text{ s.t. } \int_0^s |z_1(\tau)| d\tau \leq \Gamma(s),\ |z_1(\tau)| \leq 1\ \forall\, 0 \leq \tau \leq s \right\}. \tag{3}$$

Similarly, the cumulative demand at sale price p_2, given that the sale begins at time s, is given by:

$$N_2(T-s) = \overline{\lambda}_2\, (T-s) + \widehat{\lambda}_2 \int_s^T z_2(\tau) d\tau, \tag{4}$$

with $\widehat{\lambda}_2 < \overline{\lambda}_2$ and $z_2 \in \mathcal{Z}_2$ where $\mathcal{Z}_2$ is defined by:

$$\mathcal{Z}_2 = \left\{ z_2 \text{ s.t. } \int_s^T |z_2(\tau)| d\tau \leq \Gamma(T-s),\ |z_2(\tau)| \leq 1\ \forall\, s \leq \tau \leq T \right\}. \tag{5}$$

The robust problem is then given by:

$$\begin{aligned} \max_{0 \leq s \leq T}\ \min_{z_1 \in \mathcal{Z}_1} & \Big[p_1 \min\{K, N_1(s)\} \\ & + \min_{z_2 \in \mathcal{Z}_2} p_2 \min\{N_2(T-s), K - \min(K, N_1(s))\} \Big], \end{aligned} \tag{6}$$

with $N_1(s)$ and $N_2(T-s)$ defined by Eqs (2) and (4), respectively, and $\mathcal{Z}_1$ and $\mathcal{Z}_2$ defined by Eqs (3) and (5), respectively. Notice that the nominal problem is a special case of the robust problem, where $\Gamma(s) = 0$ for all s.

Theorem 1 (Robust Markdown Problem) *The robust problem (6) is equivalent to:*

$$\max_{0 \leq s \leq T} \left[p_1 \min\{K, N_1^-(s)\} + p_2 \min\left\{ N_2^-(T-s), K - \min\{K, N_1^-(s)\} \right\} \right], \tag{7}$$

with $N_1^-(s) = \overline{\lambda}_1 s - \widehat{\lambda}_1\, \Gamma(s)$ *and* $N_2^-(T-s) = \overline{\lambda}_2 (T-s) - \widehat{\lambda}_2\, \Gamma(T-s)$.

In particular, the worst-case revenue is achieved when demand in the first stage is lower than its nominal value.

Proof. The worst-case revenue in the second stage is always achieved when demand is less than its nominal value. (This is because the coefficient in front of $N_2(T-s)$ is positive.)

Furthermore, $p_1 \min\{K, N_1(s)\} + p_2 \min\left\{N_2^-(T-s), K - \min\{K, N_1(s)\}\right\}$ can be rewritten as:

$$(p_1 - p_2)\min\{K,\, N_1(s)\} + p_2 \min\left\{K,\, N_2^-(T-s) + \min\{K,\, N_1(s)\}\right\}$$

by adding and subtracting $p_2 \min\{K,\, N_1(s)\}$ and rearranging terms. We conclude by using that $p_2 < p_1$. □

We now characterize the optimal timing decision. We assume that $N_1^-(T) < K$, i.e., the worst-case demand at the high price is less than the total number of items initially in inventory. If this assumption is not satisfied, it is never optimal to have a sale since demand is so high that there will not be enough items to satisfy the worst-case demand at full price, and the problem becomes trivial.

Theorem 2 (Optimal Markdown Time)
Consider the equation $N_1^-(s) + N_2^-(T-s) = K$.
(i) If it does not have a solution, then it is either optimal not to have a sale, i.e., $s^* = T$ or to have an immediate sale, i.e., $s^* = 0$.
(ii) If it does have a solution, then this solution is unique and the optimal revenue is the maximum between $p_1\, N_1^-(T)$, which corresponds to the no-sale case, $p_2\, N_2^-(T)$, which corresponds to the immediate-sale case, and $p_1\, N_1^-(s^*) + p_2\, N_2^-(T - s^*)$, which corresponds to a sale starting at time s^*.

Proof. Because we have $N_1^-(T) < K$ by assumption and the worst-case demand N_1^- increases with time, we have $N_1^-(s) \leq N_1^-(T)$ and the objective function of Problem (7) can be rewritten as: $p_1\, N_1^-(s) + p_2 \min\left\{N_2^-(T-s), K - N_1^-(s)\right\}$.

For $s \in [0, T]$ such that $N_2^-(T-s) \geq K - N_1^-(s)$, the objective becomes $p_2\, K + (p_1 - p_2)\, N_1^-(s)$, which increases in s because $p_2 < p_1$ and $N_1^-(s)$ increases in s. Hence, the maximum over that set of s values is achieved when $N_2^-(T-s) = K - N_1^-(s)$ or potentially at $s = 0$, if it satisfies the condition $N_2^-(T) \geq K$. (It is easy to see that $s = T$ never satisfies the condition, and therefore is not considered in this case.)

For $s \in [0, T]$ such that $N_2^-(T-s) \leq K - N_1^-(s)$, the objective can be rewritten using the expressions of $N_1^-(s)$ and $N_2^-(T-s)$ in Theorem 1, leading to: $p_1\, [\overline{\lambda}_1 s - \widehat{\lambda}_1\, \Gamma(s)] + p_2\, [\overline{\lambda}_2 (T-s) - \widehat{\lambda}_2\, \Gamma(T-s)]$, which is a convex function because the budget of uncertainty function is concave. Hence, the maximum over that set of s values is also achieved when $N_2^-(T-s) = K - N_1^-(s)$ or potentially at $s = T$ or at $s = 0$ if $N_2^-(T) \leq K$.

We now analyze the equation $N_1^-(s) + N_2^-(T-s) = K$, i.e.,

$$(\overline{\lambda}_1 s - \widehat{\lambda}_1\, \Gamma(s)) + (\overline{\lambda}_2 (T-s) - \widehat{\lambda}_2\, \Gamma(T-s)) = K. \tag{8}$$

The left-hand-side expression is a convex function, so the equation has at most two solutions. But the extreme values of the left-hand-side expression are $N_2^-(T)$ for $s = 0$ and $N_1^-(T)$ for $s = T$. We already know that $N_1^-(T) < K$, so Eq. (8) has exactly one solution, denoted s^*, if $N_2^-(T) \geq K$, i.e., the manager can sell all items at the low price, and no solution otherwise. If no solution exists, we are left with the extremities of the season as possible sale times, and compare the revenue in the no-sale case ($s = T$ for $p_1\lambda_1 > p_2\lambda_2$) with that in the immediate-sale case ($s = 0$). If a solution to Eq. (8) does exist, we conclude by comparing the revenue at $s = 0$, $s = T$, and $s = s^*$. □

We now compare the sale times in the nominal and robust models. We say that it is optimal to put a product on sale when the sale time is strictly between 0 and T.

Corollary 3 *If it is optimal to put the items on sale both in the nominal and the robust frameworks, then the sale in the robust framework occurs earlier than in the nominal framework.*

Proof. Assume a sale before the end of the time horizon occurs in both models, at $\overline{s}$ in the nominal framework and at s^* in the robust framework. $\overline{s}$ and s^* satisfy, respectively:

$$\overline{\lambda}_1\overline{s} + \overline{\lambda}_2(T - \overline{s}) = K, \tag{9}$$

and:

$$\overline{\lambda}_1 s^* - \widehat{\lambda}_1\,\Gamma(s^*) + \overline{\lambda}_2(T - s^*) - \widehat{\lambda}_2\Gamma(T - s^*) = K. \tag{10}$$

Subtracting Eq. (9) from Eq. (10) yields:

$$(\overline{\lambda}_1 - \overline{\lambda}_2)\,(s^* - \overline{s}) - [\widehat{\lambda}_1\,\Gamma(s^*) + \widehat{\lambda}_2\Gamma(T - s^*)] = 0.$$

It follows from $\overline{\lambda}_1 < \overline{\lambda}_2$ and the positivity of the term between straight brackets that $s^* < \overline{s}$. □

The above result is particularly valuable for practitioners who might not consider putting items on sale before a certain amount of time has elapsed in the selling season. Corollary 3 shows that the decision-maker averse to demand ambiguity should be prepared to put items on sale earlier than his ambiguity-neutral counterpart.

Example 1. If the budget of uncertainty function is linear in time, i.e., $\Gamma(s) = \alpha\, s$ for some $\alpha \in (0, 1)$, we have:

$$s^* = \frac{(\overline{\lambda}_2 - \widehat{\lambda}_2\,\alpha)\,T - K}{\overline{\lambda}_2 - \widehat{\lambda}_2\,\alpha - (\overline{\lambda}_1 - \widehat{\lambda}_1\,\alpha)}, \tag{11}$$

which belongs to $[0, T]$ if and only if either:

$$\overline{\lambda}_1 - \widehat{\lambda}_1\,\alpha > \overline{\lambda}_2 - \widehat{\lambda}_2\,\alpha \text{ and } (\overline{\lambda}_2 - \widehat{\lambda}_2\,\alpha)\,T \leq K \leq (\overline{\lambda}_1 - \widehat{\lambda}_1\,\alpha)\,T,$$

or:

$$\overline{\lambda}_1 - \widehat{\lambda}_1\,\alpha < \overline{\lambda}_2 - \widehat{\lambda}_2\,\alpha \text{ and } (\overline{\lambda}_1 - \widehat{\lambda}_1\,\alpha)\,T \leq K \leq (\overline{\lambda}_2 - \widehat{\lambda}_2\,\alpha)\,T.$$

The robust approach is then equivalent to the nominal approach with the same number of items initially in inventory but *lower* arrival rates $\overline{\lambda}_1 - \widehat{\lambda}_1\,\alpha$ and $\overline{\lambda}_2 - \widehat{\lambda}_2\,\alpha$. Note that the expression of the optimal sale time in the nominal case matches that in Gallego and van Ryzin (1994).

If, in addition, the half-widths of the range forecasts are proportional to the nominal values of the arrival rates, i.e., $\widehat{\lambda}_1 = \delta\,\overline{\lambda}_1$ and $\widehat{\lambda}_2 = \delta\,\overline{\lambda}_2$ for some $\delta \in (0, 1)$, we have:

$$s^* = \frac{\overline{\lambda}_2\,T - \dfrac{K}{1 - \delta\,\alpha}}{\overline{\lambda}_2 - \overline{\lambda}_1},$$

i.e., the robust approach is then also equivalent to the nominal approach with the same arrival rates but a *higher* number of items initially in inventory. Furthermore, the higher the aversion to ambiguity, the earlier the sale (i.e., s^* decreases as α increases.)

Extension to time-varying arrival rates. In practice, arrival rates are often time-dependent (see Bitran and Mondschein (1997), Heching et al. (2002)). In that case, we will assume that the worst-case arrival rate at the sale price p_2 is always higher than any of the worst-case time-dependent arrival rate at the full price p_1. We still have to solve in s the equation $N_1^-(0, s) + N_2^-(s, T) = K$ with a slight change in notation to capture start and end times of sales rather than only durations. While the solution will depend on the structure of the arrival rate functions (as a function of time), Corollary 3 still holds. This is because we have $N_1^-(0, \overline{s}) + N_2^-(\overline{s}, T) < K$ by definition of the optimal sale time in the nominal model $\overline{s}$ and the worst-case number of arrivals N_1^- and N_2^-. Since the function $s \to N_1^-(0, s) + N_2^-(s, T)$ is decreasing in s (its slope is $\lambda_1^-(s) - \lambda_2^-(s) < 0$), we need to decrease s from $\overline{s}$ in order to have $N_1^-(0, s) + N_2^-(s, T) = K$. In other words, the optimal sale time in the robust model still occurs earlier than in the nominal model.

While some of the theoretical insights we provide in the paper require a constant-arrival-rate assumption, the mathematical models can easily be extended to the more complex case of time-dependent arrival rates by focusing directly on the worst-case number of arrivals $N_1^-(0, s)$ and $N_2^-(s, T)$.

2.1.2 Numerical Experiments

The goal of the numerical experiments is to investigate the benefits of implementing the robust framework rather than the nominal one. We will assume here that budgets are linear in time, i.e, $\Gamma(s) = \alpha s$ for some $\alpha \in (0, 1)$. This choice is motivated by the ease of the implementation of the robust optimization approach in that case.

As a simple example, consider the problem where a retailer must decide on the optimal time to mark down items from their full price at any point during a 5-month selling season. He has established that the demand at each price belongs to the intervals $[(1-\delta_1)\overline{\lambda}_1, (1+\delta_1)\overline{\lambda}_1]$ and $[(1-\delta_2)\overline{\lambda}_2, (1+\delta_2)\overline{\lambda}_2]$ respectively, with parameters summarized in Table 1.

Parameters	Full price	Sale price
(p_1, p_2)	\$10	\$9
$(\overline{\lambda}_1, \overline{\lambda}_2)$	90/month	120/month
(δ_1, δ_2)	0.2	0.2
K	500	
Optimal sale times		
$\overline{s}$	3.33 months	
$s^*(0.3)$	2.27 months	

Table 1: Problem parameters and optimal sale times in the nominal model and the robust-linear model with $\alpha = 0.3$.

In the nominal problem, it is optimal to begin the sale at $\overline{s} = 10/3$. As we have just shown, in the robust optimization approach, it is optimal to begin the sale at some time $s^*(\alpha) < \overline{s}$. For each value $\alpha = (0, 1)$ in increments of 0.1, we run 10,000 simulations for i.i.d. demands obeying a Normal, Triangular, and Uniform distributions. For example, we simulate independent Normal distribution with means $(\overline{\lambda}_1, \overline{\lambda}_2)$, and standard deviations $\left(\delta\overline{\lambda}_1/2, \delta\overline{\lambda}_2/2\right)$, respectively, with $\delta = 0.2$. We then calibrate the parameters for the Triangular and Uniform distributions such that the standard deviations are equal across all three.

Table 2 and Figures 1 and 2 illustrate the performance of our approach. Table 2 compares the performance of our nominal model versus the robust model for $\delta = 0.2$ and $\alpha = 0.3$ for each of the three distributions, with the sale beginning at time $s^*(0.3) = 2.27$. The robust model reduces the standard deviation of the realized revenues and increases the value of the 10^{th} and 25^{th} percentiles, while achieving a slightly lower mean revenue.

Figure 1 compares performance across δ for $\alpha = 0.3$ in the Normal distribution

	Normal		Triangular		Uniform	
	Nominal	Robust	Nom.	Rob.	Nom.	Rob.
mean ($)	4671	4664	4668	4663	4665	4654
25^{th} percentile ($)	4569	4682	4552	4677	4536	4662
10^{th} percentile ($)	4346	4535	4340	4515	4323	4532
st. dev. ($)	210	114	211	112	209	93

Table 2: Comparison of nominal and robust performance ($\delta = 0.2, \alpha = 0.3$).

case. As the range for the demand rates increases, the revenue distributions shift to the left and become wider, i.e., volatility and downside risk (risk that the random outcome will be *worse* than expected) increase. Figure 2 compares performance across α for $\delta = 0.2$ for the Normal distribution. As α increases, we find that standard deviation decreases and the 10^{th} percentile increases. But we also find that a risk-averse investor sometimes achieves both higher revenue and lower risk, e.g., the mean revenue and 25^{th} percentile for $\alpha = 0.1$ and $\alpha = 0.2$ are higher than the nominal values. Outside this exception, our results show a tradeoff of higher return in the nominal case and lower risk in the robust case, in particular lower downside risk. (Figures 1 and 2 exhibit similar properties for the Triangular and Uniform cases, and thus are omitted here.) The numerical simulations were performed using Excel @Risk Analysis Software.

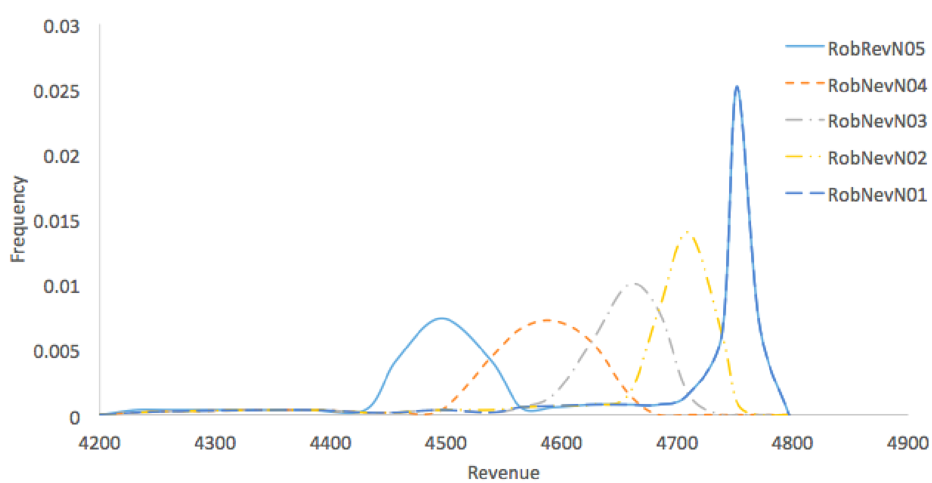

Figure 1: Revenue histogram across δ ($\alpha = 0.3$)

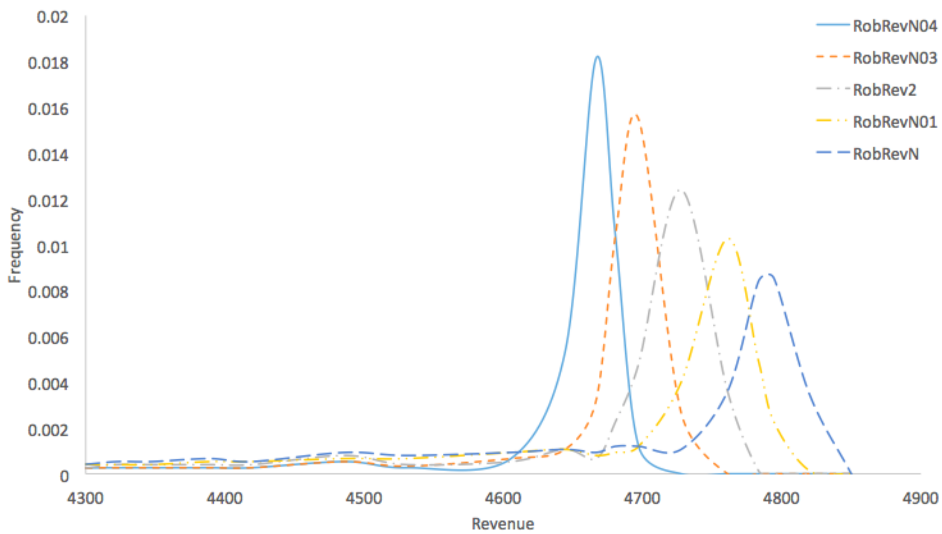

Figure 2: Revenue histogram across α ($\delta = 0.2$)

2.2 I-price Problem, I>2

We now investigate the case with multiple sale times, corresponding to multiple potential discounts. The decision-maker never increases prices during the selling season. Possible sale prices are ranked in decreasing order and considered in that order during the selling season, i.e., $p_1 > \cdots > p_I$, with p_1 the price offered when the season begins ($t = 0$). s_i is the time where the price switches from p_{i-1} to p_i. Note that there will be $I - 1$ sale times, since there are I prices. Let $\widehat{\lambda}_i = \delta_i \overline{\lambda}_i$ for some $\delta_i \in (0, 1)$ be the arrival rate of the demand when the item is priced at p_i, for all $i = 1, \dots, I$. For completeness, let $s_0 = 0$ and $s_I = T$.

Robust model. The robust optimization problem will determine how many items to offer at each price point, assuming that the demand takes its worst case value (within the uncertainty set) in each case and while respecting the inventory constraint. This is expressed mathematically as follows:

$$\begin{aligned} \max_{\mathbf{x},\mathbf{s}} \quad & \sum_{i=1}^{I} p_i x_i \\ \text{s.t.} \quad & \sum_{i=1}^{I} x_i \leq K \\ & x_i \leq \overline{\lambda}_i \left[s_i - s_{i-1} - \delta_i \Gamma(s_i - s_{i-1}) \right], \quad \forall i, \\ & x_i \geq 0, \quad \forall i. \end{aligned} \tag{12}$$

We do not need to impose constraints on the s_i since they will be implicitly enforced by the constraints on the x_i. Specifically, the implicit constraint that $s_i - s_{i-1} - \delta_i \Gamma(s_i - s_{i-1}) \geq 0$ for all i, with the function $s \to s - \delta\Gamma(s)$ increasing in s for any $0 \leq \delta \leq 1$ by assumption and taking the value 0 for $s = 0$ ensures that we will have $s_i \geq s_{i-1}$ for all i.

An important observation is that Problem (12) is linear if the budgets of uncertainty are linear in time, but non-linear and non-convex if the budgets are concave. Therefore, we will distinguish between linear and non-linear budgets to investigate the tractability of our approach.

2.2.1 Linear Budget of Uncertainty

The result below provides an important structural property of the optimal sale prices. Note that it also holds for the nominal problem, which is a special case of the robust formulation where the linear budgets are set to zero.

Theorem 4 (One Sale Time)
We always have at most one sale time when the budgets of uncertainty are linear in time, and thus always have at most two positive x_i.

Proof: By contradiction using Problem (12), which is a linear programming problem in this case. There are $2I - 1$ decision variables and $2I + 1$ constraints. Thus, $2I - 1$ constraints among $2I + 1$ must be tight in order to have a basic feasible solution, leaving two non-binding constraints. Let n be the number of x_i that are strictly positive at optimality; hence, $I - n$ of the x_i are zero (there are $I - n$ tight sign constraints), so there must be $2I - 1 - (I - n) = I + n - 1$ non-sign constraints that are tight at optimality. But there are only $I + 1$ such constraints. Therefore, we must have $I + n - 1 \leq I + 1$, i.e., $n \leq 2$. Note that if two x_i are not tight and the other $I - 2$ are tight, we have exactly one sale time; if only one x_i is strictly positive and the other $I - 1$ are zero, all inventory is sold at one price and there is either no sale or an immediate sale. □

Theorem 4 shows that the robust optimization approach with linear budgets of uncertainty reduces to a one-sale model, but the prices implemented before and after the sale must now be determined from the menu $(p_1, \ldots, p_I)$ by a linear programming problem rather than being known in advance. For instance, with a menu of three prices that allows for two markdowns, it may be optimal to go directly from p_2 to p_3 by staying for a zero amount of time at p_1, or from p_1 to p_2. This is what is determined by the LP.

In fact, we recover above part of Proposition 4 in Gallego and van Ryzin (1994), specifically, the existence of a single sale time. Gallego and van Ryzin (1994) go one step further by proving that the prices immediately follow each other, so that the two positive x_i are of the type x_j and x_{j+1} for some j. We recover that part below in Theorem 5. (In other words, with a menu of three prices that allows for two markdowns, it will never be optimal to go directly from p_1 to p_3.) This property is specific to the nominal case and the linear budget of uncertainty case. We see from the numerical example in Table 5 (for instance for $\delta = 0.4$) that this property no longer holds for concave budgets of uncertainty.

Example 2. Following Example 1, let $\Gamma(s) = \alpha\, s$ for some $\alpha \in (0, 1)$. In order to gain theoretical insights, we will make two assumptions, which we will refer to jointly as Assumption A1:

1. the δ_i parameters are constant, such that the half-widths of the range forecasts (uncertainty measures) are proportional to the nominal values of the arrival rates, $\widehat{\lambda}_i = \delta\, \overline{\lambda}_i$, for all i.

2. $i \to p_i(\overline{\lambda_i})\, \overline{\lambda}_i$ is concave (the values $(p_i(\overline{\lambda_{i+1}})\, \overline{\lambda}_{i+1} - p_i(\overline{\lambda_i})\, \overline{\lambda}_i)/(\overline{\lambda}_{i+1} - \overline{\lambda}_i)$ are non-increasing). In particular, there exists an i^* such that $(p_i(\lambda_i)\, \lambda_i)$ increases for $i \leq i^*$ and decreases for $i > i^*$. This arises for instance when the arrival rates are decreasing and linear in the prices.

We define i_0 as the greatest index i such that $\overline{\lambda}_i \leq K/[T\,(1 - \delta\,\alpha)]$. Because the $\overline{\lambda}_i$'s increase with i, $L = \{l : \overline{\lambda}_l[T\,(1 - \delta\,\alpha)] \leq K\}$ is such that $L = \{1, \ldots, i_0\}$ and $H = \{h : \overline{\lambda}_l[T\,(1 - \delta\,\alpha)] > K\}$ is such that $H = \{(i_0 + 1), \ldots, T\}$. The sets L and H are often called "abundant capacity" and "scarce capacity", respectively, in the existing literature (Bitran and Caldentey, 2003). We will assume below that all possible i^* (defined above as the maximizers of $p_i(\lambda_i)\, \lambda_i$) belong to the set H; if at least one such i^* belongs to the set L, it is optimal to sell the item at price p_{i^*} for that i^* throughout the selling horizon and the problem is trivial.

Theorem 5 (Special case) *Under Assumption A1 and $i^* \in H$:*
(i) it is optimal to first charge price p_{i_0}, and put items on sale at price p_{i_0+1} and time s^* given by:

$$s^* = \frac{\overline{\lambda}_{i_0+1}\, T - \dfrac{K}{1 - \delta\,\alpha}}{\overline{\lambda}_{i_0+1} - \overline{\lambda}_{i_0}}.$$

(ii) As the uncertainty parameter δ or the risk aversion parameter α increases, both the full price and the sale price are non-increasing and the optimal sale time decreases until it reaches 0, at which point i_0 is updated to $i_0 + 1$.

Proof: (i) This result was also proved in Gallego and van Ryzin (1994) in the nominal case; however, we provide our proof here for completeness and to bridge the nominal and the robust case. Let $1 \leq a < b \leq I$. If one sale time exists in $(0, T)$, let p_a, p_b represent the prices at which items are sold before and after the optimal sale time, respectively. We know from Example 1 that:

$$s^* = \frac{\overline{\lambda}_b T - \dfrac{K}{1 - \delta\,\alpha}}{\overline{\lambda}_b - \overline{\lambda}_a}.$$

Again, the robust approach is equivalent to the nominal approach with the same arrival rates but a *higher* number of items in inventory. The optimal objective is:

$$Obj = (1 - \delta\,\alpha)\,\left[p_a \overline{\lambda}_a s^* + p_b \overline{\lambda}_b (T - s^*)\right],$$

that is:

$$Obj = (1 - \delta\,\alpha)\,\left[p_b \overline{\lambda}_b T + (p_a \overline{\lambda}_a - p_b \overline{\lambda}_b) \frac{\overline{\lambda}_b T - \dfrac{K}{1 - \delta\,\alpha}}{\overline{\lambda}_b - \overline{\lambda}_a}\right]. \tag{13}$$

Because all i^* are in H, there is a tradeoff between being able to sell items at a slower rate $\overline{\lambda}_l$ but at a higher price p_l before the optimal sale time, and at a relatively higher rate $\overline{\lambda}_h$ but at a lower price p_h after the optimal sale time. In this case, the optimal a and b indices are such that $a \in L$ and $b \in H$.

Then maximizing Eq. (13) over $a \in \{1, \ldots, i_0\}$ yields $a = i_0$ due to the slope of concave functions being non-increasing and the coefficient in front of the slope being negative. We then maximize Eq. (13) over b. Note that Eq. (13) can be rewritten as:

$$Obj = \left[\frac{p_b \overline{\lambda}_b - p_a \overline{\lambda}_a}{\overline{\lambda}_b - \overline{\lambda}_a} K + T\,(1 - \delta\alpha) \frac{p_a - p_b}{1/\overline{\lambda}_a - 1/\overline{\lambda}_b}\right].$$

It is straightforward to prove that, if $\lambda \to p(\lambda)\,\lambda$ is concave in λ, then $1/\lambda \to p(\lambda)$ is concave as well. Therefore, the objective decreases in b and the optimal b is $b = i_0 + 1$.

(ii) follows from the fact that, as $\delta\,\alpha$ increases, i_0 is non-decreasing and $K/(1-\delta\alpha)$ increases at i_0 given. $i_0 + 1$ moves from set H to set L when $\overline{\lambda}_{i_0+1}[T\,(1 - \delta\,\alpha)] = K$, which corresponds to $s^* = 0$. □

2.2.2 Concave Budget of Uncertainty

We now investigate the case where the budget of uncertainty function is non-linear and concave in time, and does not grow faster than time ($\Gamma(s) \leq s$ for all s). The concave budget is motivated by the law of large numbers in statistics, which states, in essence, that independent random variables tend to cancel each other out when considered additively – because some will be higher than their expected value and some will be lower – so the decision-maker does not need to protect himself against all of them taking their worst-case value to maintain good risk protection. The concavity comes from the fact that this behavior is more likely to be observed as the number of independent random variables increases.

The following theorem shows that the one-sale property derived in the case of a linear budget of uncertainty is preserved in the concave case under a mild assumption. (Note however that in the concave case, price levels may no longer be consecutive, as evidenced in Table 5 of the numerical results.) For simplicity, we will derive explicitly the condition under which this property holds under the special case of the proportionality assumption: $\delta_i = \delta$ for all i. The property can also hold when the proportionality assumption is not satisfied but the condition is not particularly more insightful, especially since there is no strong case in practice to pick different δ_i parameters, and thus is omitted here.

We will assume that the $p_i\,\overline{\lambda_i}$ are concave in i and denote i^* the index that achieves the maximum of $p_i\,\overline{\lambda}_i$. Under the proportionality assumption, if $\overline{\lambda}_{i^*}\,(1 - \delta\Gamma(T)) \leq K$, it is optimal to sell all the items at price p_{i^*} over the whole selling horizon. If $\overline{\lambda}_{i^*}\,(1 - \delta\Gamma(T)) > K$, then the solution is no longer trivial and we will assume that this condition holds below. We will further assume that the revenues at each price level are "distinct enough", in the sense that:

$$\max\left(\frac{p_b\overline{\lambda}_b}{p_a\overline{\lambda}_a}, \frac{p_a\overline{\lambda}_a}{p_b\overline{\lambda}_b}\right) \geq \frac{1-\delta\,\Gamma'(T)}{1-\delta\,\Gamma'(0)}, \forall a < b.$$

This condition is rather mild; for instance, if $\Gamma'(0) = 0.47$, $\Gamma'(T) = 0.8$ and $\delta = 0.1$, the right-hand side is 1.02. It is also not tight because we bound the slope of the Γ function using the extremities of the interval. (Note that because the budget cannot grow faster than time, we must have $\Gamma'(0) \leq 1$.) This set of assumptions will be referred to as Assumption A2.

Theorem 6 (One Sale Time)
Under Assumption A2, there exists an optimal solution with at most one sale time in $(0, T)$ and therefore with at most two positive x_i.

Proof: By contradiction. Assume there are at least two indices i, denoted a and b ($a < b$), for which the number of arrivals at price p_i in the robust model is given by $\lambda_i^-(s_i - s_{i-1})$, i.e., the worst-case arrival rate times the length of the time where the product is offered at that price. This means that there must be at least a third index c where the number of sales is determined by the remaining time horizon or capacity, and so at least three indices i for which $x_i > 0$.

The part of the objective that depends on the first sale time s_a is given by:

$$(\overline{\lambda}_a s_a - \delta\overline{\lambda}_a\Gamma(s_a))p_a + (\overline{\lambda}_b(s_b - s_a) - \delta\overline{\lambda}_b\Gamma(s_b - s_a))p_b.$$

The derivative of this expression in s_a is $\overline{\lambda}_a p_a - \overline{\lambda}_b p_b - \delta[p_a\overline{\lambda}_a\Gamma'(s_a) - p_b\overline{\lambda}_b\Gamma'(s_b - s_a)]$. Because $\Gamma'(T) \leq \Gamma'(t) \leq \Gamma'(0)$, it is bounded by:

- $(1 - \delta\Gamma'(T))p_a\overline{\lambda}_a - (1 - \delta\Gamma'(0))p_b\overline{\lambda}_b$ from above
- and by $(1 - \delta\Gamma'(0))p_a\overline{\lambda}_a - (1 - \delta\Gamma'(T))p_b\overline{\lambda}_b$ from below.

Case 1: If $p_a\overline{\lambda}_a \leq p_b\overline{\lambda}_b$, we want $(1 - \delta\Gamma'(T))p_a\overline{\lambda}_a - (1 - \delta\Gamma'(0))p_b\overline{\lambda}_b \leq 0$ or

$$\frac{p_b\overline{\lambda}_b}{p_a\overline{\lambda}_a} \geq \frac{1 - \delta\,\Gamma'(T)}{1 - \delta\,\Gamma'(0)}.$$

Note that the right-hand side is greater than 1. In that case, the objective is always increased by having s_a go to 0, i.e., putting items on sale at price p_b from the start and not having items on sale at price p_a.

Case 2: If $p_a\overline{\lambda}_a > p_b\overline{\lambda}_b$, we want $(1 - \delta\Gamma'(0))p_a\overline{\lambda}_a - (1 - \delta\Gamma'(T))p_b\overline{\lambda}_b \geq 0$ or

$$\frac{p_a\overline{\lambda}_a}{p_b\overline{\lambda}_b} \geq \frac{1 - \delta\,\Gamma'(T)}{1 - \delta\,\Gamma'(0)}.$$

In that case, the objective is always increased by having s_a go to s_b, i.e., putting items on sale at price p_a from the start and not having items on sale at price p_b.

Assumption A2 formalizes the conditions of Case 1 and Case 2. □

Note that Theorem 6 only provides a sufficient condition for there to be only one sale time (two different prices) at optimality. This condition is not always satisfied in our numerical experiments, yet we do still observe one sale time at optimality.

Concave budgets of uncertainty turn Problem (12) into a non-linear, *non-convex* problem, which increases its difficulty and hinders our ability to solve this problem to global optimality. In what follows, we study a piecewise linear approximation to the concave budget function, which we will reformulate as a mixed-integer problem using binary variables. We can then solve the robust optimization problem using commercial mixed-integer programming solvers.

We first consider an inner approximation to the budget of uncertainty function, i.e., we approximate the function from below. This makes the feasible set of Problem (12) smaller and therefore allows us to obtain a *lower bound* on the optimal revenue, in line with the decision-maker's conservative attitude toward uncertainty. (The approach can also be applied to construct outer approximations to the budget curve using information on the tangent at given points, which is useful to assess the quality of the approximation.) Let J be the number of pieces of the piecewise linear approximation and J_j^{vec} its j-th breakpoint, with $J_0^{vec} = 0$ and $J_J^{vec} = T$. For $j = 0, \ldots, (J-1)$, let m_j and b_j be the slope and intercept with the y-axis of the $(j+1)$-th piece – in the inner approximation scheme, the segment connecting $(J_j^{vec}, \Gamma(J_j^{vec}))$ and $(J_{j+1}^{vec}, \Gamma(J_{j+1}^{vec}))$ – respectively.

Theorem 7 (Tractable Reformulation) *Problem (12) can be approximated by the following mixed-integer programming problem:*

$$
\begin{array}{lll}
\max\limits_{\mathbf{x},\mathbf{u},\mathbf{s}} & \sum\limits_{i=1}^{I} p_i\, x_i & \\
\textit{s.t.} & \sum\limits_{i=1}^{I} x_i \leq K, & \\
& x_i - \overline{\lambda}_i \left[(1-\delta_i m_j)(s_i - s_{i-1}) - \delta_i b_j\right] \leq K(1-u_{ij}), & \forall i, \forall j \\
& \sum\limits_{j=1}^{J} u_{ij} = 1, & \forall i, \\
& x_i \geq 0, & \forall i, \\
& u_{ij} \in \{0,1\}, & \forall i,\ j.
\end{array} \tag{14}
$$

Proof: This follows from a straightforward use of binary variables. At optimality, the u_{ij} variable will be equal to 1 if $s_i - s_{i-1}$ belongs to $[J_{j-1}^{vec}, J_j^{vec}]$ (with $j \geq 1$), in which case we must have $x_i \leq \overline{\lambda}_i \left[(1-\delta_i m_j)(s_i - s_{i-1}) - \delta_i b_j\right]$ otherwise, $\forall i, j$, and 0 otherwise. □

Extension to time-varying arrival rates. In the presence of time-varying arrival rates, the functions $N_1^-(0,s)$ and $N_2^-(s,T)$ should be approximated by piecewise linear functions using binary variables, instead of the budget of uncertainty function. The robust formulation will remain a MIP.

Observation. In the linear case, introducing uncertainty caused a risk-averse decision-maker to start the sale earlier in the selling season. In the non-linear case, however, the problem is more challenging because the manager faces two competing incentives: he can (1) put items on sale earlier, as before, or (2) expand the length of the selling period at the high price because the budget of uncertainty grows at a decreasing rate.

A disadvantage of this formulation is that it requires a large number of binary variables for realistic approximations of the budget of uncertainty function over time. On the other hand, we observe that only a few of the disjunctive constraints $x_i - \overline{\lambda}_i \left[(1-\delta_i m_j)(s_i - s_{i-1}) - \delta_i b_j\right] \leq K(1-u_{ij})$ with $\sum_{j=1}^{J} u_{ij} = 1$ will be tight at optimality. Therefore, we describe below an algorithm generating only the constraints as needed, in order to keep the problem tractable. The constraints are generated when the corresponding u_{ij} binary variable is allowed to possibly take the value 1 (instead of be kept at 0).

The preprocessing plays an important role in the speed of the algorithm in numerical experiments. We will assume that, in the nominal model, it is optimal to have a sale time within $(0, T)$, so that we have two non-zero x_i (recall that they are of the type x_a and x_{a+1}; it is straightforward to modify the preprocessing to the case where we only have x_a and not x_{a+1}. We then introduce (allow to vary) the u_{aj} and $u_{a+1,j}$ corresponding to an earlier sale time s_a in the robust model than in the nominal model, motivated by the robust-linear case. Note that if this is not true in the robust-concave model, the correct u_{ij} will then be generated later. The u_{ij} for $i \neq a$ and $i \neq a+1$ are set to 1 so that those $s_i - s_{i-1}$, previously at 0, can increase to any value in $[0, J_1^{vec}]$ and the slope of the budget function still be correctly modeled using $u_{i1} = 1$. Again, if it ends up being optimal for $s_i - s_{i-1}$ to increase beyond J_1^{vec}, the correct u_{ij} will be unfixed from 0 (allowed to vary, introduced as decision variable) later.

Algorithm 8 (Solving the MIP approximation)

1. *Step 0. (Preprocessing.) Solve the nominal model. Let a be the high price level index in the optimal solution; we know $a+1$ is then the low one.*

   - *Set $u_{i1} = 1$ and $u_{ij} = 0$ for $j \geq 2$ for all i such that $x_i = 0$.*
   - *Allow the $u_{a,j}$ to vary for all j such that $\lceil s_{0a} \rceil \geq J_j^{vec}$, set the others $u_{a,j}$ to 0.*
   - *Allow the $u_{a+1,j}$ to vary for all j such that $\lceil T - s_{0a} \rceil \leq J_j^{vec}$, set the others $u_{a+1,j}$ to 0.*

   Set $r := 1$ where r indicates the iteration number.

2. *Step r, $r \geq 1$. Solve Problem (14) over $\mathbf{x}$, $\mathbf{s}$ and only the $\mathbf{u}$ that have been unfixed. (For a u_{ij} that is still at 0, $x_i - \overline{\lambda}_i \left[(1 - \delta_i m_j)(s_i - s_{i-1}) - \delta_i b_j\right] \leq K(1 - u_{ij})$ is not generated at this step.) Obtain the new optimal solution $\mathbf{x_r}, \mathbf{s_r}, \mathbf{u_r}$.*

3. *For each i, compute $j_r(i)$ such that $s_{r,i} - s_{r,i-1}$ belongs to $[J_{j_r(i)-1}^{vec}, , J_{j_r(i)}^{vec}]$. Unfix $u_{i,j_r(i)}$ if it had not already been unfixed.*

4. *If there was no new u_{ij} to unfix from 0, stop. Else, set $r := r + 1$ and repeat.*

Optimality Bounds. The MIP model above uses an *inner* approximation to the budget of uncertainty, i.e., the piecewise linear segments fall under the curve, capturing less uncertainty in the MIP model than what the non-linear budget function would plan for. It is also possible to use the same framework for an *outer* approximation of the budget of uncertainty function, where the segments fall above the curve, and are defined by the slope of the concave budget at each point j in J^{vec}. By solving the MIP using lower and upper piecewise approximations to the budget of uncertainty, we are able to derive lower and upper bounds on the optimal objective in the robust optimization framework. Approximations with more pieces will be closer to the non-linear curve. The optimality gap can thus be narrowed down until it falls below a prespecified tolerance.

2.2.3 Numerical Experiments

The goals of these numerical experiments are twofold: (1) to compare the performance of a decision-maker using a linear and a (MIP approximation to a) concave budget of uncertainty function, and (2) to illustrate the performance of the simplification procedure where we eliminate unnecessary constraints and binary variables, leading to smaller MIPs.

We assume a 5-month selling season with units in months (the time horizon is thus $[0, 5]$), with 5 possible markdowns (10%, 20%, 30%, 40%, or 50% off) from the initial price of \$100. The parameter values are summarized in Table 3. In order to compare the MIP solution for a concave budget of uncertainty with the linear budget case, we distinguish between α^L used to define a linear budget uncertainty and α^{NL} used to define a concave budget. Then, we calibrate the parameters to reflect similar values for the solution to the nominal problem. Specifically, for a non-linear budget of uncertainty with α^{NL} and β, we find α^L such that $\alpha^L s^* = \alpha^{NL}(s^*)^\beta$.

Parameter types	Parameter values
Prices p^{vec}	$\$(100, 90, 80, 70, 60, 50)$
Nominal arrival rates per month $\overline{\lambda}^{vec}$	$(50, 70, 90, 110, 130, 150)$
Deviations from nominal values δ^{vec}	$(0.2, 0.2, 0.2, 0.2, 0.2, 0.2)$

Table 3: Example parameters.

Here, the optimal nominal solution is to sell items at $p_3 = \$80$ until $s^* = 2.5$, and then sell for the remaining time at $p_4 = \$70$, for an optimal revenue of \$37,250. We consider a concave budget of uncertainty where $\alpha^{NL} = 0.47$ and $\beta = 0.5$. We then let $\alpha^L = \frac{0.47}{\sqrt{2.5}} \approx 0.3$. Our piecewise linear approximation will have 5 pieces ($J = 5$), with each integer time period serving as a breakpoint ($J_j^{vec} = j$ for all j).

Optimal markdown times, markdown prices and corresponding revenue statistics are provided in Table 4. The revenue statistics are computed as follows. We test the optimal strategies obtained in the nominal, robust-linear, and robust-MIP models by simulating the two relevant arrival rates 2,000 times each, using independent Uniform random variables in $[(1-\delta)\,\overline{\lambda}_i,\ (1+\delta)\,\overline{\lambda}_i]$, $i = 3, 4$. Table 4 also summarizes our findings. Note that the expected revenues of the two robust models are now very close to each other, as was expected due to the small difference in optimal policies. In both the robust linear case and the MIP approximation to the concave budget, the left tail of the revenue distribution has been shifted to the right, as indicated by the higher values of the 1^{st} and 5^{th} percentiles. Although standard deviation has increased, our numerical results suggest that this is due here to an increase in upside risk, i.e., "good risk" or possibility that random outcomes will be *better* than their expected value, because the difference between the mean and the 5^{th} percentile has increased. This is an unexpected side effect of the robust optimization approach.

	Model		
	Nominal	Robust-Linear	Robust-MIP
s^*	2.50	0.90	1.03
Price before markdown	\$80	\$80	\$80
Price after markdown	\$70	\$70	\$70
Objective mean (\$)	37,271	38,046	37,983
median (\$)	37,154	37,961	37,875
1^{st} percentile (\$)	30,748	31,284	31,277
5^{th} percentile (\$)	32,177	32,345	32,352
75^{th} percentile (\$)	39,423	41,214	41,019
Standard deviation (\$)	3,049	3,707	3,619

Table 4: Comparison of simulation results for nominal, robust linear, and robust MIP performance ($\delta = 0.2, \alpha^L = 0.3, \beta = 0.5, \alpha^{NL} = 0.47$).

Additional numerical experiments (not shown here) suggest that the qualitative relationship between the three models described above holds across all values of α, β, and δ as long as these parameters are calibrated to reflect similar levels of uncertainty as described above. Therefore, besides the obvious case where the decision-maker strongly believes that the budget of uncertainty function should be strictly concave rather than linear, the piecewise linear approximation is most useful when we cannot find $\alpha^L \in (0, 1)$ using the procedure we have described. This happens when $\alpha^{NL}(s^*)^{\beta-1} > 1$, which corresponds to a very risk-averse decision-maker (high α^{NL} or high β, provided that $s^* > 1$).

To assess the practical benefits of Algorithm 8, we use the numerical setup from the experiments above and solve both the robust problem and the heuristic for δ increasing from 0 to 0.5 in steps of 0.1. We approximate the concave budget using 5 segments, connecting the integer points on the curve between the present time $t = 0$ and the end of time horizon $T = 5$. We compare performance in Table 5, where we report the number of evaluated nodes, number of linear iterations and total elapsed time for each approach. These statistics were obtained using the MINTO solver in NEOS. Algorithm 8 is very fast in this example because the optimal u_{ij} are generated during the preprocessing phase.

	Complexity						Solution
δ	Full MIP			Smaller MIP			s*
	N1	N2	N3	N1	N2	N3	
0.1	31308	138868	5.69	15	343	0.06	$s_3 = 1.75,\ s_4 - s_3 = 3.25$
0.2	31196	139845	6.11	23	397	0.05	$s_3 = 1.03,\ s_4 - s_3 = 3.97$
0.3	31125	141684	6.54	23	587	0.10	$s_2 = 0.26,\ s_4 - s_2 = 4.74$
0.4	27785	150990	6.34	23	529	0.06	$s_2 = 0.05,\ s_4 - s_2 = 4.95$
0.5	30172	165307	6.30	17	435	0.02	$s_4 = 5$

Table 5: Comparison of full MIP and smaller MIP (Algorithm 8) across δ. N1=number of evaluated nodes, N2=number of linear iterations and N4=total elapsed time (CPU seconds).

2.3 Dynamic Policy

A disadvantage of traditional robust optimization is that the optimal policy is static, i.e., it does not incorporate new information as it is realized and the problem must simply be resolved as conditions change. As argued in Bitran and Caldentey (2003), the rapid evolution of information technologies and the corresponding growth of the Internet and e-commerce make a static assumption potentially costly to a decision-maker. In many markets today, it is possible to collect valuable information (about demand, inventory levels, competitors' strategies, etc.) and process it in real time; in such settings, decision-makers should react dynamically to changes in the marketplace based on realized uncertainty over time. In the remainder of this section, we define a *policy* about the time we put items on sale based on the results of the robust optimization approach; we then compare the performance of the dynamic policy to the solution of the robust static model above.

For this approach, let $T^r(t)$ and $K^r(t)$ be the time and the inventory remaining at time t, respectively. At each time, we decide whether to continue to sell items at the current price or mark them down, as a function of $T^r(t)$ and $K^r(t)$. As time t passes, we update our model to account for realized demand. The high-level idea is to decide on a policy structure beforehand, for instance, putting items on sale when the ratio of remaining inventory to remaining time (i.e., the minimum arrival rate required to clear inventory) exceeds a threshold, and then determine the threshold based on the previous robust optimization analysis. The shape of the policy, based on remaining capacity and remaining time horizon, is motivated by its intuitive connection to the minimum arrival rate required; it is also connected to the revenue management literature such as van Ryzin and Talluri (2005), Ozer and Phillips (2012) and Bernstein et al. (2014). To clarify the setup, we focus on the following example.

Example with 2 prices. Consider the problem where a decision-maker must decide on the optimal time to mark down the price from $p_1 = \$10$ to $p_2 = \$9$ (i.e., by 10%) during a 20-week selling season. The retailer has established that the demand rates at each price point belong to the intervals $[(1-\delta_1)\overline{\lambda}_1, (1+\delta_1)\overline{\lambda}_1]$ and $[(1-\delta_2)\overline{\lambda}_2, (1+\delta_2)\overline{\lambda}_2]$ respectively, with $\overline{\lambda}_1 = 22.5$ per week, $\overline{\lambda}_2 = 30$ per week; in addition, $K = 500$ and $\delta_1 = \delta_2 = \delta = 0.2$. These numerical values are summarized in Table 6.

Parameter	Before sale	After sale
Prices (p_1, p_2)	\$10	\$9
Nominal arrival rates per week $(\overline{\lambda}_1, \overline{\lambda}_2)$	22.5	30
Deviations from nominal values (δ_1, δ_2)	0.2	0.2

Table 6: Example parameters.

In the linear budget of uncertainty case, our dynamic policy will be to mark down the price from p_1 to p_2 as soon as $\dfrac{K^r(t)}{T^r(t)}$ exceeds a threshold determined by the optimal s^* of the static problem. Since $T^r(t) = T - s^*$ and $K^r(t) = N_2^-(T - s^*)$ by definition of s^*, the threshold in this case does not depend on s^* and instead is equal to $(1 - \alpha^L\delta)\overline{\lambda}_2$.

For the concave budget of uncertainty, we use a piecewise linear approximation of $\Gamma(s) = \alpha s^\beta$, for some $\alpha, \beta \in (0, 1)$. For instance, with $J_j^{vec} = (0,\ 4,\ 10,\ 20)$, the first segment over 4 weeks is used to determine the policy $J_0^{vec} = 0 \leq s_i -$

$s_{i-1} < J_1^{vec} = 4$, which is motivated by the assumption that uncertainty cannot grow faster than time as above; the second segment is used when $4 \leq s_i - s_{i-1} < 10$, and the third segment is used when $10 \leq s_i - s_{i-1} < 20$. If the $T^r(t)$ and $K^r(t)$ cross the corresponding threshold determined by our dynamic policy, then it is optimal to markdown the price from p_1 to p_2. These numbers have been selected for illustrative purposes, but it is straightforward to make the time step as small as desired. The robust-MIP approximation is re-solved at the beginning of each week using remaining inventory and remaining time horizon to update the threshold values.

As above in the linear robust case, let $\delta = 0.2$ and $\alpha^L = 0.3$. For the MIP approximation, let $\delta = 0.2$, $\beta = 0.5$ and $\alpha^{NL} = 0.47$. The optimal sale times in the static models are summarized in Table 7.

Model	Opt. sale time
Nominal	13.33 weeks
Robust linear	9.08 weeks
Robust MIP	11.83 weeks

Table 7: Optimal sale times, switching from $p_1 = \$10$ to $p_2 = \$9$ with 500 items in initial inventory.

We wish to compare these three static solutions to our dynamic models, where we also consider the same three cases: nominal, robust-linear and robust-MIP-approximation. In the simulations, we model the true arrival of customers at p_1 and p_2 using Poisson processes with means $(1/\lambda_1,\ 1/\lambda_2)$ respectively, where λ_1 and λ_2 are generated as independent Uniform random variables with supports $[(1-\delta_1)\overline{\lambda}_1,\ (1+\delta_1)\overline{\lambda}_1]$ and $[(1-\delta_2)\overline{\lambda}_2,\ (1+\delta_2)\overline{\lambda}_2]$ respectively. Table 8 summarizes the results of our simulation for 10,000 runs using Matlab.

	Static Model			Dynamic Model		
	N	R,L	R,M	N	R,L	R,M
mean markdown times	13.33	9.08	11.83	14.18	12.60	13.11
mean revenue ($)	4621	4629	4635	4715	4716	4720
25th-percentile ($)	4458	4550	4539	4608	4621	4620
10th-percentile ($)	4187	4372	4263	4350	4411	4407
σ ($)	276	172	235	251	229	232

Table 8: Comparison of simulation results for static and dynamic models; N=nominal, R=robust, L=linear, M=MIP approximation.

Note that the first three markdown times in Table 8 (static model, cases N, R,L and R,M) are exactly equal to s^*. The next three reflect the average time in the dynamic models when the threshold that triggers the markdown is attained. We see that the static policies under-perform the dynamic policies in terms of mean realized revenues as well as in terms of 10^{th} and 25^{th} percentiles (worse average and worse left-tail). Thus, we argue that the robust dynamic approaches show the potential to be superior decision tools from the standpoint of a risk-averse decision-maker, while converting the static strategy to a dynamic policy entails very little effort.

3 Multiple Products

3.1 Generalities

The multiple product case has received much less attention in the literature, presumably because of the higher degree of complexity of multiple-product formulations (Bitran and Caldentey, 2003). In this section, we modify our earlier approach in order to solve the case with multiple products. We also describe the case where the decision-maker is limited to at most one sale time, and the case where he decides to either put a product on sale (in which case he does it at a time common for all products on sale), or not put the product on sale at all.

As before, our methodology relies on first defining the nominal problem, and then incorporating uncertainty on the arrival rates parameters using the robust optimization framework. Let N be the number of products. It is straightforward to write the nominal problem as:

$$\begin{array}{lll} \max\limits_{\mathbf{x},\mathbf{s}} & \sum\limits_{n=1}^{N}\sum\limits_{i=1}^{I} p_i^n\, x_i^n & \\ \text{s.t.} & \sum\limits_{i=1}^{I} x_i^n \leq K^n,\ \forall n, & \\ & x_i^n \leq \overline{\lambda}_i^n \left[s_i^n - s_{i-1}^n\right], & \forall i,n, \\ & x_i^n \geq 0, & \forall i,n. \end{array} \tag{15}$$

The decision-maker can further impose the s_i^n to be independent of n (so that all products are marked down at the same time, which is a reasonable assumption to make to avoid confusing customers.)

The most important feature of Problem (15) is that uncertainty on the demand rates only affects the $x_i^n \leq \overline{\lambda}_i^n \left[s_i^n - s_{i-1}^n\right]$ constraints; furthermore, each of these constraints is affected by only *one* source of uncertainty. It is then straightforward to extend the single-product formulation to the multiple-product case by using budget-of-uncertainty functions for each product, where the $x_i^n \leq \overline{\lambda}_i^n \left[s_i^n - s_{i-1}^n\right]$ constraints are replaced by:

$$x_i^n \leq \overline{\lambda}_i^n \left[s_i^n - s_{i-1}^n - \Gamma_i^n(s_i^n - s_{i-1}^n)\right]. \quad (16)$$

A disadvantage of this basic setup, which we will call the *robust-disaggregate approach*, is that it does not capture the fact that, under the assumption that demand rates are independent across products and price levels, all arrival rates are unlikely to achieve their worst-case bound (right-hand side of Eq. (16)) simultaneously. Applying a "naive" robust optimization approach where each constraint is protected against the worst case thus risks being overly conservative. The key issue for a robust optimization approach applied to multiple products is thus how to define the uncertainty set, and in particular, how to define the budget-of-uncertainty functions that parametrize the set in order to obtain formulations that combine tractability with good practical performance: should we have one function per product, or should we instead use an aggregate function?

3.2 Aggregate Uncertainty Constraints

To take advantage of the fact that the arrival rates are unlikely to reach their worst-case bound (right-hand side of Eq. (16)) simultaneously, we suggest making the following change to the problem structure: we will keep the $x_i^n \leq \overline{\lambda}_i^n \left[s_i^n - s_{i-1}^n\right]$ constraints using the nominal value of the demand rates, but we will also *aggregate* the constraints affected by uncertainty and apply robust optimization to this aggregated constraint.

The motivation is that in practice, the decision-maker will simply have left-over items if it turns out that $x_i^n > \lambda_i^n \left[s_i^n - s_{i-1}^n\right]$, where the λ_i^n are the actual unknown parameters, (or, if the manager implements a dynamic policy based on the robust optimization approach, he might simply change the timing of the next markdown). Therefore, the real-life consequences of constraint violation are rather mild; infeasibility is of little concern to the manager for good reason. It is thus more important to protect revenue in a way that incorporates conservatism while ensuring good performance when tested in simulations, which is the purpose of the constraint-aggregation technique. To the best of our knowledge, we are the first to suggest the idea of constraint aggregation in this context.

The decision-maker must determine the coefficients of the linear combination used in the aggregation process; the only restriction is that the coefficients be positive. A natural choice is to select the various product prices, since those are the coefficients of the linear combination in the objective, which is what the decision-maker wants to protect against uncertainty. We will use that choice in the remainder of the paper. Even after deciding to use prices as aggregation coefficients, the manager still faces two ways of aggregating constraints, leading to two possible budget-of-uncertainty functions:

1. Across products, for each markdown level:

$$\sum_{n=1}^{N} p_i^n x_i^n \leq \sum_{n=1}^{N} p_i^n \overline{\lambda}_i^n (s_i - s_{i-1}) - \sum_{n=1}^{N} \delta_i^n p_i^n \overline{\lambda}_i^n \int_{s_{i-1}}^{s_i} z_i^n(\tau) d\tau, \quad \forall i, \quad (17)$$

   where we have already injected that the worst case is achieved for demand rates lower than or equal to their nominal values. The budget of uncertainty function then satisfies:

$$\sum_{n=1}^{N} \int_0^t z_i^n(\tau) d\tau \leq \Gamma_i(t), \ \forall i.$$

   The worst case in the right-hand side of Eq. (17) can be computed offline and re-injected into the full robust formulation.

2. Across products and markdown levels.

$$\sum_{n=1}^{N} \sum_{i=1}^{I} p_i^n x_i^n \leq \sum_{n=1}^{N} \sum_{i=1}^{I} p_i^n \overline{\lambda}_i^n (s_i - s_{i-1}) - \sum_{n=1}^{N} \sum_{i=1}^{I} \delta^n p_i^n \overline{\lambda}_i^n \int_{s_{i-1}}^{s_i} z_i^n(\tau) d\tau, \quad (18)$$

   This allows us to define a single budget of uncertainty function across the entire time horizon. Again, we calculate the worst case offline. We will implement this second choice in the numerical study below because it most takes advantage of the "canceling-out" effect of uncertainty that makes robust optimization not overly conservative; in other words, it has the greatest potential for robust optimization to show high-quality performance in practice.

Once the worst-case problem has been solved offline in the demand rates, we substitute these expressions in the right-hand side of the equations (17) or (18) above (in our case, Eq. (18)), and solve the resulting problem as a deterministic problem. In the linear-budget case, the formulation is linear. As before, in the case of a concave budget, we introduce binary variables to use a piecewise linear approximation.

3.3 Additional Formulations

To make the mathematical formulation resemble most closely the real-life problem faced by retailers, we assume in what follows that all items must be put on sale at the same time, although the robust optimization methodology itself does not require this assumption to be tractable. This will require additional constraints on the sale times, using auxiliary binary variables. We present below the nominal formulations for (a) the problem with exactly one sale time common to all items and (b) the extension where the decision-maker can choose which items to put on sale, while the others are always sold at their initial price. (There is again at most one sale time, common to all products put on sale.) We then incorporate uncertainty using the aggregate method described in the previous section. Section 3.4 implements these models in numerical experiments.

3.3.1 Limit to one sale time

If we restrict the number of sale times to one, we must choose this unique markdown time, as well as the prices before and after markdown for each product. Note that each product n will share the same price *level* i before markdown, and i' after markdown. (It is also possible to write a formulation with exactly one markdown and different price levels for each product before and after markdown, using appropriately defined binary variables. This formulation is left to the reader.)

The nominal version of Problem (15) becomes the following MIP, with $s_0 = 0$ and $s_I = T$:

$$
\begin{array}{lll}
\max\limits_{\mathbf{x},\mathbf{s},\mathbf{y}} & \sum\limits_{n=1}^{N} \sum\limits_{i=1}^{I} p_i^n \, x_i^n & \\
\text{s.t.} & \sum\limits_{i=1}^{I} x_i^n \leq K^n, & \forall n, \\
& x_i^n \leq \overline{\lambda}_i^n (s_i - s_{i-1}), & \forall i, \ \forall n, \\
& 0 \leq s_i - s_{i-1} \leq T \, y_i, & \forall i, \\
& \displaystyle\sum_{i=1}^{I} y_i \leq 2, & \\
& y_i \in \{0, 1\}, & \forall i, \\
& x_i^n \geq 0, & \forall i, n.
\end{array}
$$

Note that $\sum_{i=1}^{I} y_i^n \leq 2$ ensures at most one sale time (two prices over the selling season) but the right-hand side can also be increased to any desired number of prices to be observed during the selling season. Removing the constraints where the y_i's appear allows us to recover the case with any number of multiple sale times, which we use for comparison in the numerical experiments in Section 3.4.

3.3.2 Extension to Sale Subsets

Here, we decide to either put a product on sale (in which case we do it at a time common for all products selected for the sale), or not put the product on sale at all. Each product n is first sold at price p_1^n. Only one sale time is allowed; however, different items can receive different markdowns. This situation occurs frequently in the retail industry: patrons are notified that a sale will begin at a certain date and given a range of possible discounts; when they visit the store, they find some items marked down, say, 25% while others are marked down 40%; others are not marked down.

We formulate the nominal problem using a set of mutually exclusive constraints, modeled using binary variables. If product n is never put on sale (a decision that we denote by $w_n = 0$ with w_n binary), then we enforce $x_1^n \leq \overline{\lambda}_1^n T$, as there is no point in having more items than the total demand; however, if it is put on sale at some point during the selling horizon ($w_n = 1$), then we enforce $x_i^n \leq \overline{\lambda}_i^n (s_i - s_{i-1})$. For the items put on sale, the decision-maker must also decide on the size of the markdown, that is, the index $i(n)$ determining the price level $p_{i(n)}^n$ for that product after markdown. The model is a MIP, with $s_0 = 0$ and $s_I = T$:

$$
\begin{array}{lll}
\max\limits_{\mathbf{x},\mathbf{s},\mathbf{y},\mathbf{w}} & \sum\limits_{n=1}^{N} \sum\limits_{i=1}^{I} p_i^n \, x_i^n & \\
\text{s.t.} & \sum\limits_{i=1}^{I} x_i^n \leq K^n, & \forall n, \\
& x_1^n \leq \overline{\lambda}_1^n T + K^n w_n, & \forall n, \\
& x_i^n \leq \overline{\lambda}_i^n (s_i - s_{i-1}) + K^n (1 - w_n), & \forall i, \ \forall n, \\
& x_i^n \leq K^n w_n, & \forall i \geq 2, \ \forall n, \\
& 0 \leq s_i - s_{i-1} \leq T \, y_i, & \forall i, \\
& \displaystyle\sum_{i=1}^{I} y_i \leq 2, & \\
& w_n \in \{0,1\}, & \forall n, \\
& y_i \in \{0,1\}, & \forall i, \\
& x_i^n \geq 0, & \forall i, n.
\end{array}
$$

We will perform constraint aggregation on the $x_i^n \leq \overline{\lambda}_i^n (s_i - s_{i-1}) + K^n (1 - w_n)$ constraints only (meaning that we keep them in the formulation with the nominal value of the parameters and apply robust optimization to a cumulative version of the constraints), and replace the uncertain parameters in $x_1^n \leq \overline{\lambda}_1^n T + K^n w_n$ by their *worst-case* values, because these constraints only matter when there is no sale at all and item n is sold at price p_1^n throughout the *whole* time horizon.

3.4 Numerical Experiments

In this section, we study how the choice of the robust optimization models impacts performance, regarding realized revenue, stability of the solution, and complexity of the model. To do so, we report the theoretical results of twelve models for the nominal, robust-aggregate and robust-disaggregate cases when there are multiple sales with and without subsets, and one sale with and without subsets. In order to compare the complexity of the various instances, we turn off the presolve functionality as well as all cuts in CPLEX 11.0 from AMPL. We report the resulting number of MIP iterations and branch and bound nodes for each instance. All instances were solved using an Intel(R) Core(TM)2 Duo CPU T7500 @ 2.20GHz processor.

Over a 5-month selling season ($T = 5$), a decision-maker can choose whether to mark down $N = 3$ items, with $I = 3$ sale times and thus 4 possible price levels. We assume the budget-of-uncertainty functions are linear. In the nominal case, all α and δ parameters are set to zero. In the robust case, $\alpha^n = 0.3,\ \forall n$, and $\delta_i^n = \delta = 0.2,\ \forall i, \forall n$. We also add, to show the versatility of the approach, the constraint that the demand arrival rates for at most $k = 2$ out of $N = 3$ items can take their worst-case values. This is modeled using additional variables in the worst-case problem that is solved offline and does not add any mathematical difficulty to the formulation. The capacity for each item K^n is set to 500 for all n. Values for the prices and arrival rates are shown in Table 9.

Prices		Nominal arrival rates	
p_i^1:	(100, 90, 65, 50),	$\overline{\lambda}_i^1$:	(100, 200, 230, 250),
p_i^2:	(100, 80, 79, 50),	$\overline{\lambda}_i^2$:	(80, 110, 230, 250),
p_i^3:	(100, 80, 79, 78),	$\overline{\lambda}_i^3$:	(80, 110, 130, 250).

Table 9: Parameter values. (Deviation from nominal arrival rates is always 0.2.)

We present our findings for the nominal case, a lower level of uncertainty, which we define as $\alpha = 0.3,\ \delta = 0.2$, and a higher level of uncertainty, which we define as $\alpha = 0.35,\ \delta = 0.4$. To test the performance of the robust optimization approach, we simulate i.i.d. demands obeying a Normal, Triangular, and Uniform distribution, but obtain similar trends in all three cases, and therefore only present the Normal case here. Tables 10 and 11 summarize the performance metrics of the three approaches (nominal, robust-aggregate and robust-disaggregate) for lower and higher levels of uncertainty, respectively. Note that multiple sale times, rather than a single one, can now be optimal. In each table, we highlight the best metric among the three approaches. For instance, for lower uncertainty with one sale time and with subsets, the highest mean revenue is achieved for the robust-aggregate model. Indeed, the robust-aggregate model shows superior performance on many (although not all) metrics. These numerical results highlight the potential of a robust-aggregate approach in revenue management, which we plan to investigate further in follow-up work.

	Subsets		No Subsets	
	# Sale Times		# Sale Times	
	One	Multiple	One	Multiple
Nominal Model				
mean	135,473	138,264	**134,401**	**136,643**
25th percentile	132,589	136,072	131,807	**135,345**
10th percentile	129,087	133,037	128,756	**133,530**
σ	4687	3785	4210	2283
Robust-Aggregate Model				
mean	**135,973**	**138,358**	134,385	135,927
25th percentile	**133,956**	**136,379**	**132,025**	133,818
10th percentile	**131,000**	133,332	129,005	131,210
σ	**3537**	3529	3954	3496
Robust-Disaggregate Model				
mean	135,763	138,116	133,838	134,570
25th percentile	133,263	136,422	131,786	134,136
10th percentile	129,910	**133,411**	**129,504**	133,365
σ	3899	**3226**	**3318**	**1603**

Table 10: Numerical Results for Lower Uncertainty ($\alpha = 0.3,\ \delta = 0.2$).

In the remainder of the paper, we omit all results for high uncertainty as they show trends similar to the case with low uncertainty; the reader is referred to Dziecichowicz (2011) for the complete set of numerical results. Table 12 shows the optimal solution and number of iterations required for the various problem types for lower uncertainty. We notice that in the cases of multiple sales with subsets and one sale without subsets, increasing the level of protection against uncertainty (from the nominal model to the robust-aggregate model to the robust-disaggregate model) makes the markdowns happen earlier. In the case of multiple sales without subsets and one sale with subsets, the markdowns in the robust-disaggregate case happen earlier than in the nominal case, but this is not true of the robust-aggregate case, where incorporating uncertainty results in (a) the first two price changes be-

	Subsets		No Subsets	
	# Sale Times		# Sale Times	
	One	Multiple	One	Multiple
Nominal Model				
mean	131,416	134,154	131,416	134,695
25th percentile	125,864	130,743	126,337	132,005
10th percentile	118,848	123,092	119,594	127,981
σ	9302	7772	8522	4867
Robust-Aggregate Model				
mean	**132,795**	**134,572**	**132,795**	**134,717**
25th percentile	**128,832**	**131,266**	128,672	**132,065**
10th percentile	**122,652**	**125,021**	123,719	128,073
σ	**7144**	6531	6634	4767
Robust-Disaggregate Model				
mean	131,755	133,996	132,711	130,124
25th percentile	127,659	130,937	**128,687**	129,623
10th percentile	121,711	124,800	**123,917**	**128,193**
σ	7457	**6242**	**6541**	**1907**

Table 11: Numerical Results for Higher Uncertainty ($\alpha = 0.35,\ \delta = 0.4$).

ing delayed in the case of multiple sales without subsets, and (b) the first two price changes occurring even before those in the robust-disaggregate model in the case of one sale with subsets. Interestingly, here, the choice of products the manager puts on sale (when he can select them) does not depend on the model implemented: he always selects items 2 and 3.

Figure 3 shows the histograms of the revenues for the nominal, robust-aggregate and robust-disaggregate models with lower levels of uncertainty when multiple sale times are allowed and there are no subsets (all products are put on sale). As expected – because of the conservative way we have incorporated uncertainty in the disaggregate case – the robust-disaggregate model leads to more conservative solutions; its histogram is much narrower than the other two and shifted to the left (low revenues). Figure 4 is the counterpart of Figure 3 when only one sale is allowed. The robust-aggregate model offers again the best performance.

Implementing a dynamic approach based on the results of the robust optimization model in this context can only be done in a straightforward manner if the products are considered independently of each other, allowing separate sale times. In practice, this is not a realistic behavior for the retailer, and the question of im-

Model	Obj ($)	s*	Sale Items	Iterations
Multiple sales without subsets				
1) Nominal	138,300	$[0, 3.25, 4.125, 4.75, 5]$	All	LP
2) Robust-Aggregate	137,026	$[0, 4.15, 4.15, 4.52, 5]$	All	LP
3) Robust-Disaggregate	134,625	$[0, 2.54, 3.93, 4.69, 5]$	All	LP
Multiple sales with subsets				(MIP, BB)
1) Nominal	141,728	$[0, 3.81, 3.81, 4.66, 5]$	[2,3]	(19, 3)
2) Robust-Aggregate	141,323	$[0, 3.69, 3.69, 4.58, 5]$	[2,3]	(53, 5)
3) Robust-Disaggregate	136,743	$[0, 3.43, 3.43, 4.55, 5]$	[2,3]	(36, 2)
One sale without subsets				(MIP, BB)
1) Nominal	136,471	$[0, 4.41, 4.41, 4.41, 5]$	All	(25, 6)
2) Robust-Aggregate	135,909	$[0, 4.36, 4.36, 4.36, 5]$	All	(22, 3)
3) Robust-Disaggregate	131,724	$[0, 4.22, 4.22, 4.22, 5]$	All	(25, 6)
One sale with subsets				(MIP, BB)
1) Nominal	139,412	$[0, 4.41, 4.41, 4.41, 5]$	[2,3]	(30, 6)
2) Robust-Aggregate	139,080	$[0, 3, 3, 5, 5]$	[2,3]	(72, 11)
3) Robust-Disaggregate	133,870	$[0, 4.22, 4.22, 4.22, 5]$	[2,3]	(57, 13)

Table 12: Comparison of solutions and solution times for $\alpha = 0.3, \delta = 0.2$ (lower level of uncertainty).

plementing dynamic policies for multiple products is left to future work.

4 Conclusions

We have presented a robust optimization approach to markdowns in the retail industry under uncertain demand arrival rates. To prevent over-conservatism, we have used budget-of-uncertainty functions, which depend on the time the items remain on sale at a given price point; for linear budgets, the problem structure remains linear, while we have developed piecewise linear approximations in the non-linear concave case to solve the resulting non-convex problem as a MIP. We have also developed theoretical insights into the optimal markdown times. We have further shown that a dynamic policy whose parameters are determined by the solution of the robust optimization problem is a valuable decision tool that exhibits superior performance to the static approach in numerical experiments. In the case of multiple products, we have proposed the use of constraint aggregation to improve the performance of the robust optimization methodology; constraints with

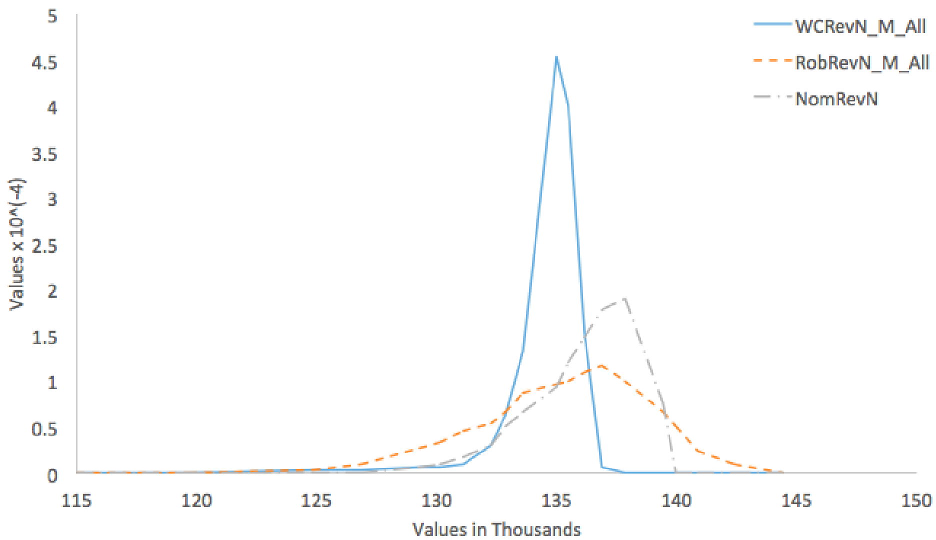

Figure 3: Multi-sale, No subset - Revenue histograms for nominal, robust-aggregate, and robust-disaggregate models ($\alpha = 0.3,\ \delta = 0.2$)

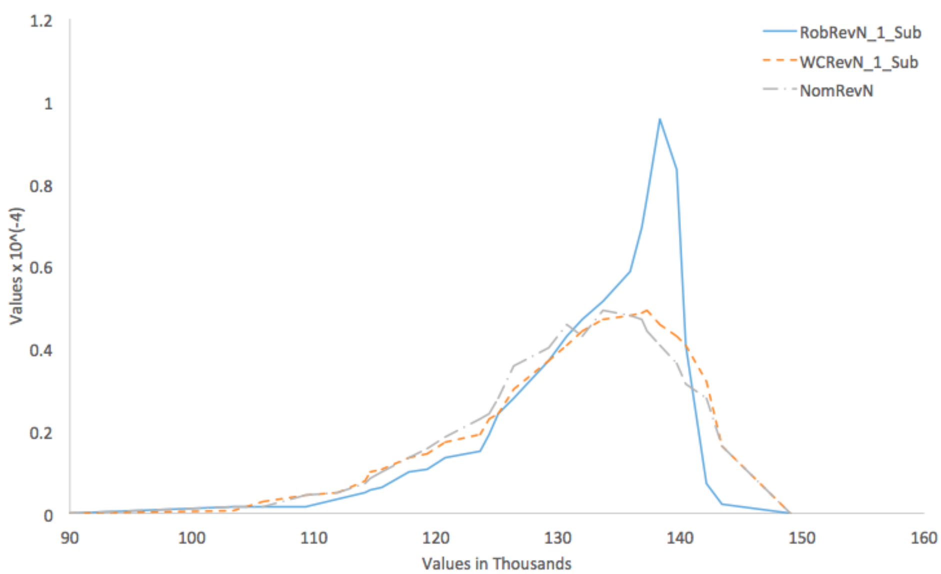

Figure 4: One-sale, Subset - Revenue histogram for nominal, robust, and worst-case model ($\alpha = 0.3,\ \delta = 0.2$)

only one uncertain parameter are solved for the nominal value of that parameter, and robust optimization is applied to a cumulative (aggregate) version of these constraints. This is possible here because, in practice, infeasibility does not occur for the type of problems considered: the decision-maker is simply left with extra products, or runs out of items. The performance of the robust-aggregate framework is very encouraging.

Future work will study methods of extending the dynamic approach to the multiple-products case, and their tractability. We would also like to study robust markups rather than markdowns; such problems arise for instance in airline revenue management.

Acknowledgements

This work was performed while both the first and second authors were students at Lehigh University. All three authors would like to thank an anonymous reviewer for his/her particularly insightful comments on connections with existing literature and suggestions to improve readability. The first author gratefully acknowledges the support of a NSF IGERT fellowship. The second author gratefully acknowledges the support of a Lehigh University Presidential Fellowship. This work was supported in part by the third author's IBM Faculty Award and NSF Grant CMMI-0757983.

References

E. Adida and G. Perakis. A robust optimization approach to dynamic pricing and inventory control with no backorders. *Mathematical Programming*, 107:97–129, 2006.

F. Bernstein, A. Gurkan Kok, and L. Xie. Dynamic assortment customization with limited inventories. Technical report, Duke University, 2014.

D. Bertsimas and G. Perakis. *Mathematical and Computational Models for Congestion Charging*, volume 101 of *Applied Optimization*, chapter Dynamic Pricing: A Learning Approach, pages 45–79. Springer, 2006.

D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1): 35–53, 2004.

D. Bertsimas, D. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

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

G. Bitran and S. Mondschein. Periodic pricing of seasonal products in retailing. *Management Science*, 43:64–79, 1997.

M. Dziecichowicz. *Robust timing decisions with applications in finance and revenue management*. PhD thesis, Lehigh University, 2011.

W. Elmaghraby and P. Keskinocak. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science*, 49(10):1287–1309, 2003.

Y. Feng and G. Gallego. Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science*, 41(8):1371–1391, 1995.

V. Gabrel, C. Murat, and A. Thiele. Recent advances in robust optimization: an overview. *European Journal of Operational Research*, 235(3):471–483, 2014.

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

A. Heching, G. Gallego, and G. van Ryzin. An empirical analysis of policies and revenue potential at one apparel retailer. *Journal of Revenue and Pricing Management*, 1(2):139–160, 2002.

O. Ozer and R. Phillips, editors. *The Oxford Handbook of Pricing Management*. Oxford University Press, Oxford, UK, 2012.

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

G. van Ryzin and K. Talluri. An introduction to revenue management. *INFORMS TutORials in Operations Research*, pages 142–194, 2005.

T. Whitin. Inventory control and price theory. *Management Science*, 2:61–68, 1955.