COLUMBIA UNIVERSITY
Center for Pricing and Revenue Management

WORKING PAPER SERIES: NO. 2013-2

Revenue Maximization via Sequential Bilateral Negotiations

Nur Ayvaz-Cavdaroglu,
Kadir Has University

Soulaymane Kachani,
Columbia University

Costis Maglaras,
Columbia University

2013

http://www.cprm.columbia.edu

Revenue Maximization via Sequential Bilateral Negotiations

Nur Ayvaz-Cavdaroglu* Soulaymane Kachani† Costis Maglaras‡

April 4, 2013

*Kadir Has University, Department of Business Administration, Istanbul, TURKEY 34083. (nur.cavdaroglu@khas.edu.tr)

†Columbia University IE&OR department, Rm. 313, School of Engineering and Applied Sciences Columbia University 500 West 120th Street New York, NY 10027. (kachani@columbia.edu)

‡Columbia Business School, 409 Uris Hall, 3022 Broadway, NY, NY 10027. (c.maglaras@gsb.columbia.edu)

Abstract

We study a dynamic pricing problem for a monopolist seller that operates in a setting where each potential sale takes the form of a bilateral negotiation. The outcome of each negotiation between the seller and a single independent buyer depends on the valuations of the seller and buyer for that good, their relative negotiation power, and their beliefs regarding the other party's valuation. We first analyze variations of the bilateral negotiation problem and analyze the effect of the buyer's negotiation power. Next, we review the dynamic negotiation problem, and propose a simple deterministic "fluid" analogue. The main emphasis of the paper is in expanding the above formulation to the case where both the buyer and seller have limited prior information on their counterparty valuation. Our first result shows that if both the seller and buyer are bidding so as to minimize their maximum regret, then it is optimal for them to bid as if the unknown valuation distributions were uniform. Building on this result and the fluid formulation of the dynamic negotiation problem, we characterize the seller's minimum price at any given point in time. Finally, we expand on the above ideas to study the seller's problem in the case where the primitives of the buyer valuation distributions are unknown and non-stationary using ideas from scenario-based robust optimization. The motivating application is from residential real-estate, however, the model and proposed approach is generally applicable.

Keywords: bilateral negotiations, dynamic pricing, revenue maximization

1 Introduction

Many transactions between a seller and a buyer follow some form of a negotiation. This is typical in business-to-business settings as well as in transactions that involve end consumers for expensive items such as cars, furniture, and real-estate. The outcome of each such negotiation depends on the reservation values of the seller and buyer, their negotiation skills, and their beliefs about these parameters for their respective counterparties. This process is known as "bilateral price negotiation". Depending on the market conditions, the seller may enjoy increased market power and as such be able to name her list price, whereas in the other extreme the buyers may essentially submit take-it-or-leave-it bids to the seller. In most settings, actual behavior falls somewhere in between, where the seller and buyer somehow split the difference between the seller's minimum acceptable bid (her reservation price) and the buyer's willingness to pay. This might be regarded as if the seller foregoing some of her profits by offering the buyer a price discount. However, in today's business world, "the shifting balance of power has many stores scrambling for pricing strategies that get beyond the time-worn cycle of markups and discounts" [9]. That is, in the new business world, there is a significant and permanent power shift towards the consumers' end, propelled by the Internet and apps, which has rendered the buyers more empowered in haggles, thus demanding much lower prices.

One motivating application for the paper comes from the residential real estate industry, where a developer of a multi-unit project, e.g., a multi-story condominium development, tries to sell various condos to prospective buyers through a sequence of negotiations over time. While for each buyer their respective negotiation could be modeled as a one-off interaction, the seller should consider the fact that she will engage into a sequence of such negotiations over time. The phenomenon of "power shift towards the buyers' end" is also observed in the real estate industry since the explosion of the real-estate bubble in the financial crisis of 2007. Hence, our main focus is the changing trading problem and the new pricing strategies of the sellers in real estate, even though most of our findings do apply to the general case. [1]

In more detail, we study the revenue maximization problem of a vendor that has C units of capacity to sell over a time horizon of length T to a market of prospective buyers that arrive according to a Poisson process with rate Λ, each has a willingness-to-pay that is an independent draw from a distribution F_b, and who engages in a bilateral negotiation with the seller for a single unit. The salvage value of the seller is private information, and buyers assume that it follows some distribution F_s and is constant over time. The reservation price of the seller at time t depends on the salvage value and the state of the sales process, i.e., the time-to-go and remaining capacity.

The ultimate focus of this paper is to study this problem, primarily in the setting where buyers have market power (which is regarded as a "buyer posted price" (BPP) environment), and where the seller and the buyers do not know the distributions F_b, F_s, respectively and moreover the unknown distribution F_b may be changing over time. This setting is motivated by the real estate application, where there is significant uncertainty about the current and future market conditions, and the

[1] Numerous articles in the press exemplify the phenomenon that "properties once sold at very high monetary terms are now being purchased by the bidders who pay the minimum amount to cover back taxes, interest and fees" [29]. Many developers of multi-unit residential projects are advertising in the newspapers, magazines and on the internet announcing that all bids are welcome.

non-stationarity here is not due to seasonality effects that can be readily incorporated, but rather due to changes in underlying market conditions, e.g., such as interest rates, economic conditions, etc., that "modulate" the buyer willingness-to-pay distribution.

Despite the importance and prevalence of negotiation problems in practice, most literature in quantitative dynamic pricing has focused on posted price mechanisms (see Gallego and Van Ryzin [13], Das Varma and Vettas [30]) and auctions (see Vulcano et al. [31]). Among the papers that involve revenue management problems in the form of bilateral negotiations, the work of Bhandari and Secomandi [7] is perhaps closest to ours regarding the problem under consideration. However, Bhandari and Secomandi use a stylized MDP to investigate the negotiation processes in a dynamic deterministic setting while we are mainly interested in uncertain environments. Moreover, our focus is not on the mechanism design, nor does it involve "strategic buyers" who refuse to buy at high prices, which are the main differences of our work from Riley and Zeckhauser [28] and Gallien [14].

The first modeling and methodological contribution of the paper is in formulating the classical bilateral negotiation problem in an uncertain environment, where buyers and the seller do not have information about F_s, F_b, respectively. There are three natural ways to specify this type of model uncertainty. The first one is stochastic, wherein the unknown distributions are assumed to be drawn from a given set of possible distributions according to some known probability law, and where the firm's goal is to optimize its expected revenues over all possible market model realizations. Its main shortcoming is that it requires detailed information on the distribution of the model uncertainty. As a second formulation, both the seller and the buyer adopt a max-min criterion where they aim to optimize their respective worst-case revenues. This criterion may yield overly pessimistic results. Finally, a third approach that reduces the conservatism of max-min formulations while maintaining their appealing low informational requirements is through the use of the competitive ratio or maximum regret criteria, which measure the performance relative to that of a fully-informed decision maker. They have been used extensively in the computer science literature, and have recently been applied in pricing and operations management problems. Specifically, Ball and Queyranne [4], Eren and Maglaras [11], Perakis and Roels [27], Lan et al. [20] and Eren and Van Ryzin [12] adopt different versions of this idea.

Secondly, we carry the analysis to the dynamic setting. The key finding is to recognize that in the buyer's market (i.e. BPP setting) where the seller is simply making accept or reject decisions of the buyer bids, the problem can be reduced to a single resource capacity control problem in the form analyzed by Lee and Hersh [21]. Specifically, the distribution of buyer bids is analogous to a continuous distribution of fare classes. This observation allows us to completely characterize the structure of the optimal policy. We note in passing that the problem in the seller's market is similarly analogous to the well-studied dynamic pricing problem in Gallego and van Ryzin [13].

Next, motivated by the goal of studying the dynamic settings where the distributional assumptions may not be known, we start with a simpler approximated problem where the buyer arrival process is replaced by a deterministic and continuous process. This model can be justified as a limit as the capacity and market potential grow large and the sales horizon and distributional assumptions stay unchanged. This is often referred to as a "fluid" model and admits a static solution, as it could be expected from the mapping of the BPP formulation to the capacity control problem, where the seller accepts all bids above a given threshold.

Finally, the last part of the paper focuses on the real-life applications where the distributions F_s, F_b are unknown and may vary over time. Motivated by our previous findings regarding the static uncertain problem, we propose a method that a) uses the fluid model, b) adopts uniform distributions for F_s, F_b, c) considers multiple possible parameter scenarios for the evolution of these distributions, and d) picks a feedback pricing strategy for the seller to optimize its regret relative to the full information problem. This problem can be solved in an open-loop manner. This, however, can be improved by optimizing over a set of linear feedback bidding rules for the seller, that are motivated by the optimal seller strategy under full information. A set of numerical results show that the regret formulation and the associated uniform distribution assumption lead to good results, i.e., modest revenue loss for the seller, in a variety of settings.

The main contributions of the paper are as follows: First, the maximum regret formulation and associated results are novel, and important on their own right as they offer a robust analogue of the one-to-one bilateral negotiations problem. Parenthetically, we find that the uniform distribution appears as the natural assumption under incomplete information, which is consistent with results derived in the robust optimization literature. Secondly, we draw attention to the analogy between the dynamic bilateral negotiation problems and the classical revenue management problems; which is a first in the literature. Third, the formulation of the seller's dynamic problem with uncertain F_s, F_b distributions assumed as being uniform, as motivated by the result in the one-to-one setting, is novel and the formulation is itself readily solvable producing a simple and tractable policy that has a good performance.

The remainder of the paper. In section 2, we consider the one-to-one negotiation problems: In section 2.1, the classical models are revisited; and in section 2.2, we analyze a variant of the problem with an added uncertainty element in terms of the valuation distribution functions. In section 3, the analysis is carried to a dynamic setting. Section 3.1 sheds light on the analogy of the negotiation and the revenue management problems. Section 3.2 presents the dynamic pricing model that extends the results of the static negotiation problem to a dynamic setting using a fluid model approach. Next, in section 4, the results of section 2.2 are extended to the dynamic setting under a regret criterion. In particular, we propose a scenario-based robust optimization approach which is both tractable and takes into account the unfolding uncertainty in the system as time progresses. Numerical illustrations and extensions are presented in section 5. Finally, section 6 concludes our findings and presents avenues for further research.

2 1-to-1 Bilateral Negotiation Problem

2.1 Background: 1-to-1 Bilateral Negotiation Problem

The literature of two-person bargaining games goes back to Nash [26] and Harsanyi [18], and the ones to pioneer the analysis of the dynamics of an environment with shifting negotiation power are Myerson (et al.) ([25], [24]) and Chatterjee and Samuelson [8]. In these studies the problem is analyzed within a static context as a game between a single seller and a single buyer.

The one-to-one bilateral negotiation problem involves the trading interactions between two individuals where one of the individuals (the seller) owns an object that the other (the buyer) wants to buy. Both players are risk neutral. From the seller's perspective the valuation of the buyer for this unit is random variable v_b, distributed according to probability density and distribution functions f_b and F_b with support $[\underline{v}_b, \bar{v}_b]$. A symmetric argument holds for the buyer, where he assumes that the seller's valuation for the unit, v_s, is distributed according to cumulative distribution function

F_s (with pdf f_s) on the range $[\underline{v}_s, \bar{v}_s]$. F_s and F_b are both strictly increasing and differentiable on their supports, and are common knowledge.

The rules of the bargaining game is as follows: At the beginning of the sales interval the seller sets a reservation price $s(v_s)$, then the buyer submits a bid $b(v_b)$, and a successful trade is concluded if $b(v_b)$ exceeds $s(v_s)$. The resulting sales price is $kb(v_b) + (1 - k)s(v_s)$, where $k \in [0, 1]$ is a parameter that determines the bargaining power of the buyers. Specifically, if $k = 0$, the problem reduces to a "seller posted price" (SPP) setting where the trade is concluded at the price $s(v_s)$ as long as $s(v_s) \leq b(v_b)$. At the other extreme $k = 1$, the problem becomes a "buyer posted price" (BPP) formulation where the sales price is equivalent to the buyer's bid $b(v_b)$, again provided that $s(v_s) \leq b(v_b)$ holds. Chatterjee and Samuelson [8] characterize the class of equilibria for this problem in which player bidding strategies are "well-behaved". In particular, they make the following assumption regarding the buyer and seller bidding functions $s(.)$ and $b(.)$, which is also relevant for our analyses:

Assumption 1. *In the equilibrium, both $b(.)$ and $s(.)$ are bounded above and below and are strictly increasing and differentiable except possibly at the boundary points.*

Under the above assumption, the equilibrium bidding strategies of the two parties [2] are the solutions to the following two linked differential equations:

$$-kF_s(s^{-1}[b(v_b)])s'(s^{-1}[b(v_b)]) + f_s(s^{-1}[b(v_b)])(v_b - b(v_b)) = 0, \tag{1}$$

$$(1-k)(1 - F_b(b^{-1}[s(v_s)]))b'(b^{-1}[s(v_s)]) + f_b(b^{-1}[s(v_s)])(v_s - s(v_s)) = 0, \tag{2}$$

where $k \in [0, 1]$ is the parameter determining the bargaining power of the buyer.

The above formulations is obtained by solving the following "best response problems" of the seller and the buyer simultaneously:

$$\max_{s \in [v_s, \bar{b}]} \int_s^{\bar{b}} (kb + (1-k)s - v_s)g_b(b)db,$$

and

$$\max_{b \in [\underline{s}, v_b]} \int_{\underline{s}}^{b} (v_b - kb - (1-k)s)g_s(s)ds,$$

[2]We will use the terms "bidding function" and "bidding strategy" interchangeably throughout the paper.

where g_s and g_b are the pdf's of the optimal bidding functions $s^*(.)$ and $b^*(.)$ respectively, $\underline{s}$ is the minimum value the seller's bid can take and $\bar{b}$ is the maximum value the buyer's bid can assume.

The equations (1) and (2) take the following simpler forms in the BPP environment (i.e. $k = 1$):

$$b^*(v_b) = \{b| - F_s(b) + (v_b - b)f_s(b) = 0\}, \quad \forall v_b \in [\underline{v}_b, \bar{v}_b] \tag{3}$$

$$s^*(v_s) = v_s, \quad \forall v_s \in [\underline{v}_s, \bar{v}_s] \tag{4}$$

and the same equations produce the following bidding functions in the SPP ($k = 0$) case:

$$b^*(v_b) = v_b, \quad \forall v_b \in [\underline{v}_b, \bar{v}_b] \tag{5}$$

$$s^*(v_s) = \{s|1 - F_b(s) + f_b(s)(v_s - s) = 0\}, \quad \forall v_s \in [\underline{v}_s, \bar{v}_s] \tag{6}$$

An interesting feature of the seller's optimal bidding function in the BPP setting is its independence from F_b. The intuition behind this fact is obvious: Since the seller has no influence on determining the final price, she is willing to accept any offer above her own valuation to obtain positive return. That makes bidding her own valuation, v_s, her best response to all bids of the buyers. Thus, g_s becoomes identical to f_s in the BPP setting and the buyer bidding function assumes the simple form as in (3). A symmetrical argument holds for the SPP setting, justifying (5) and (6).

2.2 1-to-1 Bilateral Negotiation Problem in Uncertain Environments

In this subsection, we analyze a variant of the classical one-to-one bilateral negotiation problem with an added uncertainty feature. In particular, we assume that both agents are able to estimate the minimum and the maximum values that their opponent's valuation could assume; however, they do not have any knowledge regarding the distribution of this value in its given range.

As discussed in Section 1, there are various ways to model this type of uncertainty, and among those, we will adopt the "absolute regret minimization criterion" approach (ARMC). The rationale behind this method is to improve the average quality of decisions under uncertainty.

Adopting the ARMC approach, the problems that the seller and the buyer need to solve in order to minimize their maximum regret are formulated respectively as follows:

$$\operatorname*{argmin}_{s}\{\max_{b}\max_{s'}[(kb+(1-k)s'-v_s)\cdot 1_{\{b\geq s'\}}-(kb+(1-k)s-v_s)\cdot 1_{\{b\geq s\}}]\} \tag{7}$$

$$\operatorname*{argmin}_{b}\{\max_{s}\max_{b'}[(v_b-(kb'+(1-k)s))\cdot 1_{\{b'\geq s\}}-(v_b-(kb+(1-k)s))\cdot 1_{\{b\geq s\}}]\} \tag{8}$$

In the first of the above problems, the seller tries to select the bid s which minimizes the revenue loss across all bids b of the buyer; where the seller's revenue loss in each instance is the difference between the maximum revenue she could have achieved by bidding her best response s' (i.e. $(kb+(1-k)s'-v_s)\cdot 1_{\{b\geq s'\}}$) and the realized revenue at her selected bid s (i.e. $(kb+(1-k)s-v_s)\cdot 1_{\{b\geq s\}}$). The problem of the buyer is symmetrical.

The equilibrium bidding functions s^*_{ARMC} and b^*_{ARMC} that solve the above problems and are best responses to each other are characterized in the following theorem.

Theorem 1 (**Equivalence of ARMC and the uniform distribution case**). *When each party in the bilateral negotiation game only possesses the support information of the opponent's value distribution and uses ARMC to maximize revenues, the equilibrium bidding functions are given as:*

$$s^*_{ARMC}(v_s)=\frac{v_s}{2-k}+\frac{(1-k)\bar{v}_b}{2}+\frac{k(1-k)\underline{v}_s}{2(2-k)},\ \forall v_s\in[\underline{v}_s,\bar{v}_s], \tag{9}$$

$$b^*_{ARMC}(v_b)=\frac{v_b}{1+k}+\frac{k\underline{v}_s}{2}+\frac{k(1-k)\bar{v}_b}{2(1+k)},\ \forall v_b\in[\underline{v}_b,\bar{v}_b]. \tag{10}$$

*which are also the equilibrium bidding functions of a game where F_s, F_b are both **uniform** on the given ranges.*

For the proof, please refer to Appendix 7.1.

The above result brings a theoretical motivation to use uniform distribution as the opponent's distribution function when there is no information. In other words, the results of the ARMC analysis support the intuition that the valuation of the counterparty could be anywhere over its support with equal probabilities when nothing is known regarding its distribution.

Remark 1. *We have thus extended the literature on the one-to-one negotiation problem where neither the seller nor the buyer know each other's distribution function, but they both know the* *range of the opponent valuations. It is also possible to analyze a third case where one of the parties is informed about the other's value distribution function, while the other only knows the range of his opponent's valuation. The analysis regarding this case can be found in Appendix 7.2.*

3 Dynamic Bilateral Negotiation Problems

3.1 The Analogy of Revenue Management and Bilateral Negotiation Problems

We next turn our attention to the main motivating problem of this paper: The revenue maximization problem of a firm that has C units to sell over a time horizon of length T to a market of prospective buyers that arrive according to a Poisson process with rate Λ, each has a willingness-to-pay that is an independent draw from a distribution F_b, and each engages in a bilateral negotiation with the seller for one unit of that good. The salvage value of the seller is private information, and buyers assume that it is drawn from some distribution F_s, and is constant over time.

The key observation in the dynamic setting is that the buyers in the system are "naive": they ignore the competition with other buyers in the market, and bid according to the equilibrium bidding function $b^*(\cdot)$ characterized by the equation (1). However, the seller will engage into a sequence of such negotiations over time, therefore submits her bid with the objective of maximizing her *overall* revenues. Thus, the seller's bid is no longer determined by the equation (2).

First, consider the BPP (buyer posted price) setting: In this setting, given the arrival rate Λ_t and the buyer bidding function b^*_{BPP}, it is possible to define the expected "sales rate" at instant t as $N(t) = N(s_t) = \Lambda_t \bar{G}_b(s_t)$. Then, the seller's revenue maximization problem takes the form:

$$\max_{\{s_t, t=1,\ldots T\}} \quad \mathbf{E}_{\xi, b_t}[\sum_{t=1}^{T}(b_t - v_s)\xi(t; N)] \tag{11}$$

$$\text{subject to} \quad \sum_{t=1}^{T}\xi(t; N) \leq C \; a.s., \; s_t \in \mathbb{P}, \; \forall t. \tag{12}$$

where $b_t := b^*_{BPP}(v_t)$, and the valuation v_t of the buyer arriving at t is randomly drawn from the distribution of the buyer values. $\xi(t; N)$ is the random sales amount at t which is Bernoulli with probability $\mathbb{P}(\xi(t; N) = 1) = N(t)\delta t$ and $\mathbb{P}(\xi(t; N) = 0) = 1 - N(t)\delta t$ for small δt. Note that the price s_t (which belongs to a feasible price set $\mathbb{P}$) has no direct effect on the revenue, except for determining the lower bound of the buyer bids to be admitted. That is, it effectively works as a control that leads to "opening" product classes (buyer bids) that exceed s_t and "closing" classes that bring lower revenue than s_t.

Hence, the problem is in the same spirit as the "capacity control problem" of the revenue management literature, which is studied by Lee and Hersh [21] among many others. In this problem the prices are exogenously determined by competition or through a higher order optimization problem defining the market conditions and the firm chooses a dynamic capacity allocation rule. To see the connection more clearly, assume that we approximate all buyer bids by n finite values; i.e. define $\bar{b}^* \geq b_1 \geq b_2 \geq \ldots \geq b_n \geq \underline{b}^*$ as n finite "fare classes", where the arrival rate of bid b_i is approximated by $\Lambda_t(\bar{G}_b(b_i) - \bar{G}_b(b_{i-1}))$, $\forall i \in \{2, \ldots, n\}$ and the arrival rate of b_1 is approximated by $\Lambda_t \bar{G}_b(b_1)$ at each instant t. Then, the problem above pours into the following capacity allocation problem of a firm which has discretion as to which product requests to accept at any given time:

$$\begin{aligned} \max_{\{u(t), t=1,\ldots T\}} \quad & \mathbf{E}_\xi[\sum_{t=1}^{T}(b' - v_s)\xi(t; u\Lambda)] \\ \text{subject to} \quad & \sum_{t=1}^{T} e'\xi(t; u\Lambda) \leq C \; a.s.,\; u_i(t) \in \{0,1\},\; \forall t. \end{aligned} \tag{13}$$

where $u_i(t)$'s are the controls that take value 1 when a bid of value b_i is accepted at time t and zero otherwise, $b' = \{b_1, b_2, \ldots, b_n\}$, e' the n-dimensional unit vector, and $\xi(t; u\Lambda)$ denotes the associated sales vector. The formulation (13) is the discretized version of the capacity control problem of Lee and Hersh [21]. Thus, if the buyer bids could be approximated by a finite class of fares, the BPP formulation is equivalent to the capacity allocation problem of a seller selling a single resource to multiple demand classes in a perfect competition setting. For details, we refer the reader to Lee and Hersh [21] and Maglaras and Meissner [23].

Next consider the SPP (seller posted price) setting: In this environment, the seller's revenue maximization problem could be formulated as follows:

$$\max_{\{s_t, t=1,\dots T\}} \quad \mathbf{E}_{\xi}[\sum_{t=1}^{T}(s_t - v_s)\xi(t;N)] \tag{14}$$

$$\text{subject to} \quad \sum_{t=1}^{T}\xi(t;N) \leq C\, a.s.,\; s_t \in \mathbb{P},\; \forall t. \tag{15}$$

where $N(t) = N(s_t) = \Lambda_t \bar{F}_b(s_t)$ is the "sales rate" at t and $\xi(t;N)$ is the associated sales vector. As in the BPP setting, a simple observation shows that the above formulation is in parallel to the "dynamic pricing" problem of a monopolist seller selling a homogenous product in a discrete-time setting; which is readily given in the paper of Gallego and Van Ryzin [13]. The stochastic dynamic pricing game has been extensively analyzed in the same paper, which we will be referring to as GVR paper in the sequel.

The two above equivalences stem from the fact that while each buyer negotiates with the seller only once, the seller will engage into a sequence of negotiations over the sales horizon. Hence, in BPP, she will determine the minimum bid to be accepted at each instant to control *the amount of capacity to be sold*, whereas in SPP she will pursue a *dynamic pricing strategy* to maximize the revenues to be extracted from the stochastically arriving buyers. Therefore, in broad terms, the SPP setting reduces to the dynamic pricing problem and the BPP setting to the capacity allocation problem of the literature. We state this result as a proposition.

Proposition 1. *If the buyers in the market are naive, the dynamic SPP game becomes equivalent to the* ***dynamic pricing problem*** *and the dynamic BPP game to the* ***capacity allocation problem*** *of the revenue management literature.*

3.2 Fluid Formulation of the Dynamic Problem

Since analyzing the stochastic dynamic pricing problem of the seller is difficult, we will proceed with a fluid formulation hoping to obtain insights towards the solution of the stochastic problem. As commonly known, fluid formulation is a good approximation of the real stochastic problem when number of interactions per unit time is sufficiently large.

To this end, consider the following fluid version of the dynamic negotiation game: Infinitesimal buyers arrive with a (deterministic) rate Λ_t at t, $t \in [0, T]$. Both parties know Λ_t and the distribution function of their opponent. Then the revenue maximization problem of the seller is as follows:

$$\max_{s_t, \forall t} \quad [\int_0^T r_t(v_s, s_t)dt] \tag{16}$$

$$\text{subject to} \quad \int_0^T \Lambda_t [\int_{s_t}^{\bar{b}} g_b(b)\, db]\, dt \leq C \tag{17}$$

where $r_t(v_s,\, s_t)$ is the instantaneous net revenue function of the seller at time t when her valuation is v_s and her reservation price s_t; which is given by:

$$r_t(v_s, s_t) = \int_{s_t}^{\bar{b}} \Lambda_t (kb + (1-k)s_t - v_s)\, g_b(b)\, db. \tag{18}$$

and g_b is the pdf of the buyer bidding function $b(\cdot)$ characterized in (1) and $s(\cdot)$ is given by (2).

If the above problem is modeled as a stochastic control problem in the price space, finding its solution could be extremely difficult. Therefore, following a similar approach as in GVR, we will analyze the problem by focusing on the optimal *sales rate*, rather than the optimal *pricing policy.*

If the seller sets s_t as the lowest price to be accepted at t, the fraction of buyers that are accepted at that instant is given by $\alpha_t(s_t) = \int_{s_t}^{\bar{b}} g_b(b)\, db = \bar{G}_b(s_t)$, inducing an inverse function:

$$s_t(\alpha_t) = G_b^{-1}(1 - \alpha_t).$$

The function $s_t(\alpha_t)$ is well-defined for all $\alpha_t \in [0, 1]$ as a result of Assumption 1.

Then, the instantaneous net revenue function of the seller at time t in terms of the fraction of accepted buyers becomes:

$$r_{t,a}(v_s, \alpha_t) = \int_{G_b^{-1}(1-\alpha_t)}^{\bar{b}} \Lambda_t \left(kb + (1-k)(G_b^{-1}(1-\alpha_t)) - v_s\right) g_b(b)\, db.$$

Thus, the seller's revenue maximization problem (16)-(17) in the price space is equivalent to:

$$\max_{\alpha_t, \forall t} \quad [\int_0^T r_{t,a}(v_s, \alpha_t)dt] \tag{19}$$

$$\text{subject to} \quad \int_0^T \Lambda_t \alpha_t\, dt \leq C. \tag{20}$$

which is a formulation in the demand space.

Provided that $r_{t,a}(v_s,\, \alpha)$ is concave in α, the formulation (19)-(20) becomes maximization of a concave function over a convex set; and its solution is then given as in the following Theorem.

Theorem 2. *If $r_{t,a}(v_s,\, \alpha)$ is concave in α, the equilibrium bidding strategy $s_t(.)$, $t \in [0,T]$, of the seller in the dynamic negotiation problem takes the form:*

$$s_t(v_s) = \max\{G_b^{-1}(1 - \frac{C}{\int_{t=0}^{T} \Lambda_t \, dt}), s^*(v_s)\}, \forall t \in [0,T], \tag{21}$$

where $s^*(v_s)$ is the bidding function of the seller given in (2); and the equilibrium bidding strategy $b_t(.)$ of each infinitesimal buyer arriving at time t is characterized by (1), with G_b being its cdf.

Proof. As we have already noted, the buyers have neither the knowledge of the sales rate nor the remaining inventories of the seller. Therefore, they regard the situation as a one-to-one negotiation game and employ the equilibrium bidding function $b^*(.)$ regardless of their arrival time.

To see how the seller behaves, note that the problem (19)-(20) is maximized at the maximizer of $r_{t,a}(., v_s)$, which is $\alpha^* := \bar{G}_b(s^*(v_s))$, provided that it is feasible to admit this fraction at each instant t (i.e. if $\alpha^* \int_0^T \Lambda_t \, dt \leq C$). This case is equivalent to applying the bid $s_t(v_s) = s^*(v_s)$, $\forall t$.

If, on the other hand, $\alpha^* \int_0^T \Lambda_t \, dt > C$, then by the concavity of $r_{t,a}(., v_s)$, it is optimal to admit the constant fraction $\alpha_0 := \frac{C}{\int_{t=0}^{T} \Lambda_t dt}$ at each t. This second case corresponds to bidding $s_t(v_s) = G_b^{-1}(1 - \frac{C}{\int_{t=0}^{T} \Lambda_t \, dt})$, $\forall t \in [0,T]$. So the seller will set her reservation price as $s_t(v_s) = \max\{s^*(v_s), G_b^{-1}(1 - \frac{C}{\int_{t=0}^{T} \Lambda_t \, dt})\}$, which ends the proof of the theorem. $\square$

The above theorem is in the same spirit as the Proposition 2 of GVR paper and forms the first major result of this section.

Regarding concavity of the instantaneous revenue function of the seller, for instance:

$$g_b'(b) \geq 0, \forall b \in [\underline{b}, \bar{b}] \tag{22}$$

is a sufficient condition to ensure that $r_{t,a}(v_s,\, \alpha)$ is concave in α for all $v_s \in [\underline{v}_s, \bar{v}_s]$. This condition simply ensures that the second derivative of the function $r_{t,a}(v_s, \cdot)$ is negative at all α. Observe that if both functions F_s, F_b are uniform, Condition (22) is satisfied.

3.3 Dynamic Negotiation Problems under Uncertainty

In this part, we study a variant of dynamic negotiation problems where the primitives of the buyer valuation distribution are unknown.

The problem setting is as follows: At each instant t, $t \in [0, T]$, independent negotiations take place between the seller and the entire population of infinitesimal buyers whose valuation distribution function is revealed at t. The players know each other's distribution range for all $t \in [0, T]$ (and suppose that, for convenience, this range does not change across time). In this situation, the ARMC (absolute regret minimization) approach is again a viable choice for all parties. However, we need to make the following assumption to ignore the "learning effect" for the seller (otherwise, the seller's problem becomes trivial as she can infer the value distribution function of buyers from the instantaneous sales rate and employ the optimal pricing policy).

Assumption 2. *The seller can neither observe the buyer value distribution function, F_b, nor the sales amount until the end of the sales horizon.*

Although the above assumption might seem unrealistic, it is in fact equivalent to assuming that the buyers' valuation distribution is continuously changing over time. Hence, observing the past sales will not help the seller in predicting the future sales.

With these observations, we are ready to state and prove the following Theorem, which emphasizes the analogy of the dynamic stochastic problem with the stochastic one-to-one problem.

Theorem 3. *The dynamic stochastic problem with unknown valuations reduces to the dynamic deterministic problem of section 3.2, with F_s and F_b being uniform distribution functions on their given ranges at each t.*

The proof of the above Theorem can be found in the Appendix 7.3.

4 Applications in Non-Stationary Environments

In this section, we consider dynamic stochastic problems where the support of valuation distributions are unknown and non-stationary. This type of multi-stage stochastic optimization problems has elicited much interest from various research communities and there are several established methodologies to expound them involving dynamic programming, stochastic programming and robust optimization. However, the problem usually remains hard to solve analytically. Therefore, in practice, it is typical to solve the recursions numerically or resort to some approximations such as approximate dynamic programming or simulation. In a similar manner, we will introduce a class of policies that is motivated by the structure identified in the deterministic version of the problem and confirm that these policies achieve "good" performance in the dynamic stochastic problem.

Before proceeding with the analysis, first we would like to shed light on the relevance of the results of Lan et al. [20] and Lobel and Perakis [22] to our problem, where both papers analyze the capacity rationing problem of a seller operating under limited demand information. Both papers employ a robust formulation approach and resort to "absolute regret minimization criterion" (ARMC) among others. The resulting optimal policies are in the form of a nested booking policy. However, as pointed out earlier, despite the analogy between the dynamic BPP (buyer posted price) problem and the classical capacity rationing problem, restricting the buyer bids to a fixed set of discrete fares and characterizing the worst-case scenario by a specific sequence of buyer arrivals (as in these papers) would only be analyzing a special case of the stochastic BPP problem. We will rather proceed with the general problem where we allow for a continuous range of buyer bids changing dynamically over time, and assume no specific sequence or volume of buyer arrivals.

The problem with added time-varying nature of the valuations could seem to be far-fetched to the reader; however, it is commonly observed in some business settings, particularly in the real-estate sector. The only caveat is to assume that the parameters of the distribution fluctuate over night could be unrealistic. To avoid this, we will use approximations such that these transitions are significantly observable only from one discrete period to another, where the length of each period could be different from another. The following example describes such a setting:

Example: Consider a condo-developer who has $C = 375$ units to sell over $T = 15$ bi-monthly intervals. Assume that the market conditions remain stationary within an interval, but there is an observable transition in the buyer valuation distribution at the end of each period. In particular, the buyer valuations in period t, $t \in \{1, 2, \ldots, 15\}$, are uniformly distributed in the range $[\mu(t) - \$300K, \mu(t) + \$300K]$, where $\mu(t)$ is given by the equation:

$$\mu(t) = \mu(t-1) + \delta(t), \text{ for } t = 2, 3, \ldots T; \quad \mu(1) = \$600K$$

and $\delta(t)$ is the noise factor at time t with the following distribution:

$$\delta(t) = \begin{cases} \mathbf{U}[-\$120K, 0], & \text{w.p } 0.4 \\ \$0, & \text{w.p } 0.2 \\ \mathbf{U}[0, \$120K] & \text{w.p } 0.4 \end{cases} \quad \forall t$$

That is, $\mu(t)$ corresponds to average buyer valuation in period t. Some example μ paths (i.e. the values that $\mu(.)$ takes from $t = 1$ to $t = 15$) are given in the following Figure.

–Figure 1 "Example of Parameter paths" about here–

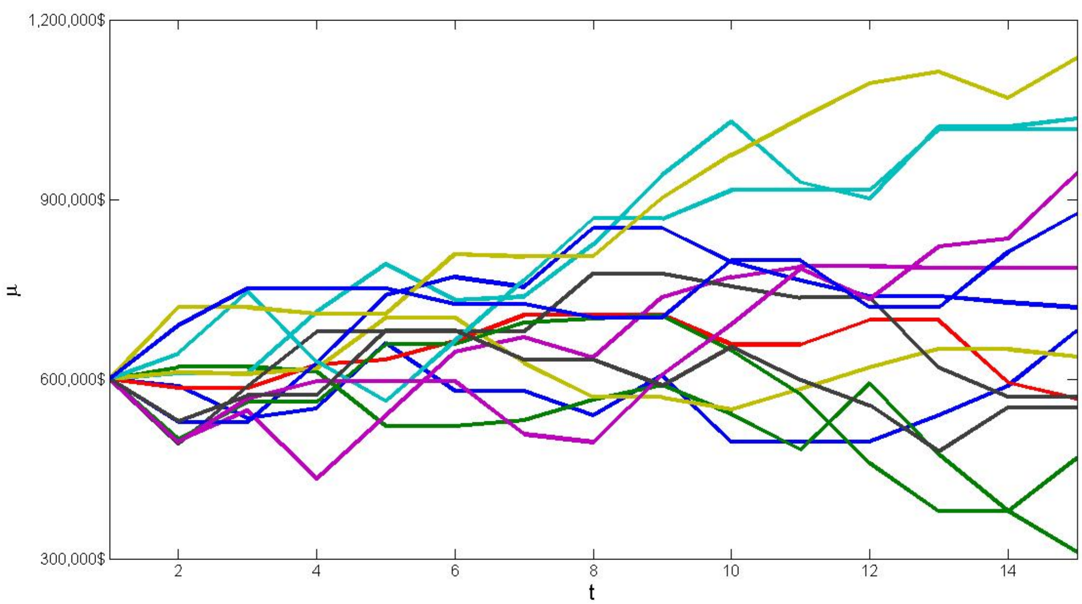

Figure 1: Examples of Parameter Paths

As exemplified above, the problem setting we will consider is comprised of discrete periods $t \in \{1, 2, ...T\}$. In each period, the buyer valuations follow a new distribution function (denoted by $F_{\mu(t)}$), and independent dynamic negotiations take place between the seller and the entire population of infinitesimal buyers. The seller regards $F_{\mu(t)}$ as uniform distribution over an unknown range (the reasoning for this assumption comes from Theorem 3). To parameterize the uniform distribution with a single variable, we will assume that the length of the distribution support is known and given at all times t; but the middle point of the range, $\mu(t)$, remains unknown. This average buyer valuation changes from one period to another by an additive noise factor. We make no distributional assumptions regarding the noise factor, except that it lies in a basic compact algebraic set $\mathbf{\Delta}$. We will focus on a BPP (buyer posted price) setting, in the light of the previous discussion about our motivating problem. For simplicity, we will pursue the analysis on the example problem stated above.

Resizing the problem by dividing all monetary values by \$300K and carrying the analysis to a fluid setting, we obtain the following robust optimization problem:

$$\max_{\{s_t, t=1,...15\}} \quad z \tag{23}$$

$$\text{subject to} \quad z \leq [\sum_{t=1}^{15} \frac{100}{4}((1.5 + \sum_{i=1}^{t} \delta(i))^2 - 4s_t^2)], \tag{24}$$

$$100((1.5 + \sum_{i=1}^{t} \delta(i)) - 2s_t) \leq a_t, \quad \forall t, \tag{25}$$

$$\sum_{t=1}^{15} a_t \leq 375 \quad a.s. \, \forall t \tag{26}$$

$$-0.2 \leq \delta(t) \leq 0.2, \quad a_t \geq 0, \, \forall t. \tag{27}$$

The formulation above is that of an uncertain quadratically constrained (QC) problem. This class of problems is analyzed by many researchers, including Ben-Tal et al. [1], who build an SDP which approximates the NP-hard robust counterpart, and Goldfarb and Iyengar [16] who reformulate it as an SOCP problem and solve the latter. Although the solution methodologies in these papers decrease the computational effort considerably, the main problem is that the above formulation leads to an open-loop solution (i.e. a pricing policy s_t that does not make use of the past disturbances), therefore yielding conservative results for practical use.

Rather than open-loop policies that do not take into account the system dynamics, some simple but tractable functional forms might be sufficient for good performances, if not for optimality. "An affine policy" of past disturbances, i.e. a pricing policy of the form: $s_t = m_t + \sum_{i=1}^{t} B_{t,i}\delta(i)$, could be one such policy. This approach is not new in the literature. It has been originally advocated in the context of stochastic programming (see Garstka and Wets [15] and references therein), where such policies are known as decision rules. More recently, the idea has received renewed interest in robust optimization (Ben-Tal et al. [2]), and has been extended to various contexts for solving specific optimization problems, from linear and quadratic programs (Ben-Tal et al. [3], Kerrigan and Maciejowski [19]) to conic and semi-definite (Ben-Tal et al. [3], Bertsimas and Brown [5]).

Bertsimas et al. [10] were able to show that the optimal controls in a one-dimensional, discrete, linear, time-varying dynamical system are affine functions of the past disturbances. However, they also show that the optimality of affine policies is easily violated even when the problem assumptions are relaxed slightly. For instance, optimality no longer holds when there exist linear constraints coupling the controls across different time-steps, which is exactly the case in our problem due to the capacity constraints. Hence, in our problem affine policies are not optimal, but hopefully yield results that are close to optimal.

Moreover, we have an additional structural support in the favor of using affine policies. To see this, let us take a step back and consider the deterministic problem. The form of the optimal pricing policy in the deterministic problem is given in the following Proposition.

Proposition 2. *If the entire μ path is known at $t = 1$ (i.e. the noise vector δ is known), the optimal bidding function of the seller with valuation v_s engaging in dynamic negotiations under BPP setting is given by:*

$$\hat{s}_t(v_s) = max\{b^*(v_b^0), v_s\}, \forall t \in \{1, 2, \ldots, T\}$$

where v_b^0 satisfies: $\sum_{t=1}^{T} \Lambda_t \int_{v_b^0}^{\mu(t)+0.5l(t)} f_{\mu(t)}(v_b)dv_b = C$ ($l(.)$ being the range of buyer valuations at t, $\Lambda(t)$ total number of buyers in the market at t).

Hence, the optimal clairvoyant policy of the seller (under ARMC approach) is a stationary policy where the optimal bid at any time t is given by: $s_t^*(v_s) = \max\{v_s, \frac{\sum_{t'=t}^{T} A_{t'}(\frac{\mu(t')+0.5l(t')}{l(t')}) - x(t)}{2\sum_{t'=t}^{T}\frac{A_{t'}}{l(t')}} + \frac{v_s}{2}\}$, where $x(t)$ is the inventory at the beginning of period t.

Thus, inspired by the optimal policy of the deterministic problem, a candidate closed-loop policy for the stochastic problem could be defined as:

$$s_t = A_t + B_t x(t) + C_{t,t}\mu(t) + \sum_{i>t} C_{i,t}(\hat{\mu}(i) + \mathbf{E}[\delta(i)]) \tag{28}$$

for appropriate constants A_t, B_t and $C_{i,t}$, $\forall t$, $\forall i > t$.

However, optimizing over the coefficients A_t, B_t and $C_{i,t}$ violates the convex nature of the maximum regret minimization problem (23)-(27), since $x(t)$ is dependent on s_i, $i = 1, 2, \ldots, t-1$, $\forall t$. Fortunately both $x(t)$ and $\mu(t)$ are functions of the noise factors $\delta(i), i \leq t$; and it is possible to recover the form (28) by defining the optimal policy s as an affine function of the past uncertainties:

Proposition 3. *Defining:*

$$s_t = m_t + \sum_{j=1}^{t} B_{t,j}\delta(j) \tag{29}$$

the formulation (28) can be recovered.

The proof of the above proposition is omitted for length-related concerns, but it simply follows from the fact that, if the seller bid is defined as in (29), the current capacity $x(t)$ is an affine function of previous noise factors $\delta(j)$, $j \leq t$, in a uniform distribution setting. However, the opposite of this claim is not true, i.e. it is not possible to recover equations (29) from the (28). That is because the degree of freedom is larger for the set of equations (29) (i.e. given the values of m_t and $B_{t,i}$, $\forall t, i \leq t$, there is more than one solution for A_t, B_t, $C_{i,t}$, $\forall t, i > t$.)

Hence, supported by previous research and the structural form of the deterministic optimal policy, we confine our search to affine pricing policies. Moreover, rather than accounting for the entire uncertainty set, we will sample N (to be found by trial-and-error) scenarios and model the seller's pricing problem with the objective of "minimizing the worst case regret" within this sample. This approach avoids computational complexity and is supported by previous works. For instance, Perakis and Roels [27] argue that rather than spanning the entire uncertainty set, accounting for the twenty-fifth and seventy-fifth percentiles produces policies that perform substantially better on average without deteriorating much in terms of the worst-case regret performance.

To this end, we define the following quantities:

$\Pi(s,\delta) :=$ net revenues to be obtained by the pricing policy s under the noise vector δ;

$\Pi^*(\delta) :=$ maximum revenues to be obtained under the noise vector δ.

Clearly, for the example problem (14)-(15):

$$\begin{aligned} \Pi(s,\delta) &= \frac{100}{4}\Big[\sum_{t=1}^{15} 0.25((1.5+\sum_{i=1}^{t}\delta(i))^2 - 4(s_t)^2)\Big] \\ \Pi^*(\delta) &= \frac{100}{4}\Big[\sum_{t=1}^{15} 0.25((1.5+\sum_{i=1}^{t}\delta(i))^2 - 4(s_t^*(\delta))^2)\Big] \end{aligned}$$

where $s_t^*(\delta) = \max\{v_s, \frac{\left[\sum_{t=1}^{15} 100\times(1.5+\sum_{i=1}^{t}\delta(i))\right]-375}{2\sum_{t=1}^{15} 100}\}$, $\forall t$.

And the final form of the problem to be solved is the following:

$$\begin{aligned} \min_{\{m_t, B_{t,i}, t=1,\dots 15, i=1,\dots,t\}} \quad & z \\ \text{subject to} \quad & z \geq [\Pi^*(\delta^j) - \Pi(s^j,\delta^j)], \quad \forall j = 1,2,\dots,N, \\ & 100((1.5+\sum_{i=1}^{t}\delta^j(i)) - 2s_t^j) \leq a_t^j, \quad \forall t, \forall j = 1,2,\dots,N, \\ & \sum_{t=1}^{15} a_t^j \leq 375 \quad a.s. \, \forall j = 1,2,\dots,N, \\ & s_t^j = m_t + \sum_{i=1}^{t} B_{t,i}\delta^j(t), \quad \forall t, \forall j = 1,2,\dots,N, \\ & a_t^j \geq 0, \quad \forall t, \forall j = 1,2,\dots,N. \end{aligned}$$

Hence, by finding the best common coefficients $m_t, B_{t,i}$ that minimize maximum regret across the selected scenarios, we hope to find a heuristic policy that also performs well for all possible instances of the problem.

5 Numerical Results

5.1 The Effect of the Negotiation Parameter

In this subsection we analyze the effect of the "buyer's negotiation power" (which is reflected in the parameter k) on the seller revenues. To this end, we consider a dynamic setting where the buyers and the seller both have uniform valuation distributions on the ranges $[\underline{v}_b, \bar{v}_b] = [1, 3]$ and $[\underline{v}_s, \bar{v}_s] = [0.5, 1.5]$ respectively. Assume that the buyers arrive according to a Poisson distribution with rate $\Lambda = 1$ per period for a sales horizon of $T = 50$ periods. Recall that the buyer and the seller bidding functions in the dynamic problem for a given value of the parameter k take the forms:

$$b^*(v_b) = \frac{v_b}{1+k} + \frac{k\underline{v}_s}{2} + \frac{k(1-k)\bar{v}_b}{2(1+k)},\ \forall v_b \in [\underline{v}_b, \bar{v}_b] \tag{30}$$

$$s_t^*(v_s) = \max\{v_s, G_b^{-1}(1 - \frac{x(t)}{\int_{\tau=t}^{T} \Lambda_\tau d\tau})\},\ \forall t, \forall v_s \in [\underline{v}_s, \bar{v}_s] \tag{31}$$

respectively, where $x(t)$ is the remaining inventory at t, and $G_b(.)$ is the cdf of $b^*(.)$.

We vary the value of k from 0 (i.e. SPP setting) to 1 (i.e. BPP setting) and use 500 random instances. The ratio of average seller revenues for the given k value to the revenues under the SPP setting at various levels of seller capacity is given in the Figure 2.

–Figure 2 "Seller revenues (as % of revenue at $k = 0$) for various k and C values" about here–

Although an SPP environment essentially yields higher profits for the seller than a BPP setting as expected, an interesting observation is that the seller with high load factor might actually benefit from a slight shift in negotiation power. This is because, the buyer bids might first increase and then decrease in k for lower-valued buyers. (For instance, take a buyer with $v_b = 1.2$. His bid will be equivalent to $b(v_b) = v_b = 1.2$ for $k = 0$; $b(v_b) = \frac{v_b}{1.2} + \frac{0.2\times0.5}{2} + \frac{0.2\times0.8\times3}{2\times1.2} = 1.25$ for $k = 0.2$, and $b(v_b) = \frac{v_b}{2} + \frac{0.5}{2} = 0.85$ for $k = 1$.) As the load factor $\frac{C\Lambda}{T}$ increases, it is more prevalent to accept lower-valued buyers, who now bid highest at moderate values of k rather than at $k = 0$ or at $k = 1$.

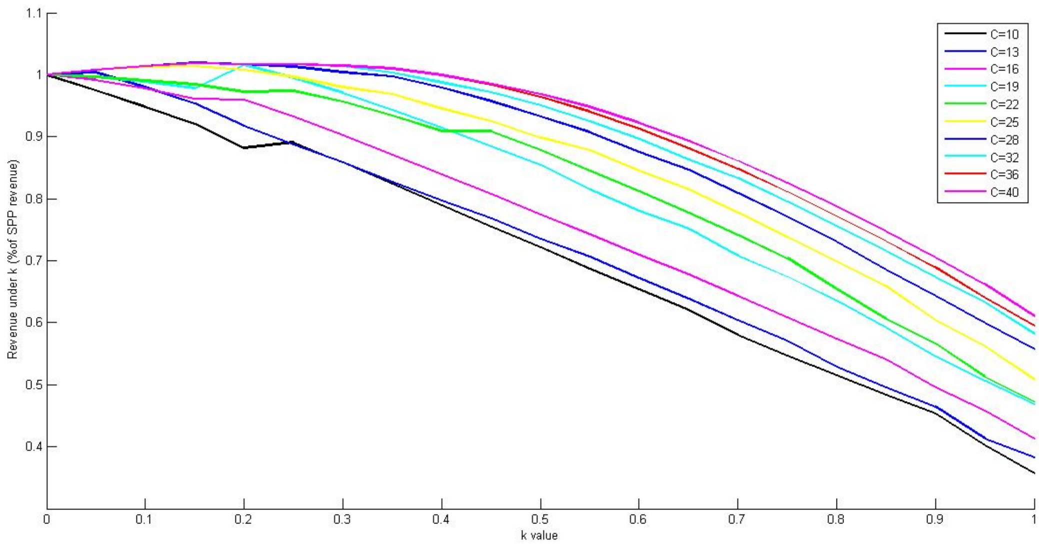

Figure 2: Seller revenues (as % of revenue at $k = 0$) for various k and C values

5.2 The Effect of Uniform Distribution Assumption

Next, we would like to investigate the seller's loss when she does not have the real distribution information and assumes that the buyers' valuations are distributed uniformly in their range as a natural conclusion of the ARMC approach. Our experiments contrast the revenues obtained by the seller in the "no distribution information" setting to the revenues in the "full-information" setting. To this end, consider the revenue maximization problem of a seller who operates in a BPP setting, where the market size is Poisson with rate Λ=100 per period for $T = 15$ periods.

For the Normal and Gumbel distributions, we extracted the mean as the midpoint of the range and selected the standard deviation σ by assuming that the range is equal to $\pm 3\sigma$. For the exponential distribution we assumed that the valuation of a typical consumer is given by $\underline{v}_b + w$ where w is exponentially distributed in $[0,\, \bar{v}_b - \underline{v}_b]$ and its rate parameter μ is selected so that the probability that w lies in that range is 99.5% (this is consistent with the $\pm 3\sigma$ assumption of the Normal distribution). In each test case, we assumed that the buyers bid believing that the seller's value is uniform in $[\underline{v}_s,\, \bar{v}_s] = [\$750\text{K},\, \$2000\text{K}]$; inducing $b^*(v_b) = \min\{v_b, 0.5v_b + 0.5\underline{v}_s\}$.

The sets of results summarized in Tables 1-2 illustrate the performance of the policy under uniform distribution assumption in a variety of settings as we varied the range ($[\underline{v}_b,\ \bar{v}_b]$), the inventory of the seller (C), and the seller valuation(v_s). In Table 1, v_s is fixed at v_s = \$1000K, while C and [$\underline{v}_b$, $\bar{v}_b$] are varied to test different cases. In Table 2, C is fixed at C = 500 where [$\underline{v}_b$, $\bar{v}_b$] and v_s are varied. We display the revenues of the no-information case as a percentage of the revenues of the full-information case (i.e. maximum revenues to be achieved).

"Table 1 about here"

Table 1: The Ratio of Seller's Revenue under ARMC to Seller's Revenue under Full Information

	[$\underline{v}_b$, $\bar{v}_b$] = [\$500K, \$1500K] C = 250, C = 500, C = 750	[$\underline{v}_b$, $\bar{v}_b$] = [\$1000K, \$2000K] C = 250, C = 500, C = 750	[$\underline{v}_b$, $\bar{v}_b$] = [\$1000K, \$2500K] C = 250, C = 500, C = 750	[$\underline{v}_b$, $\bar{v}_b$] = [\$1000K, \$3000K] C = 250, C = 500, C = 750
Exponential	96.49%, 100%, 100%	44.83%, 71.19%, 90.65%	45.60%, 64.43%, 83.99%	47.13%, 63.95%, 80.34%
Normal	92.65%, 100%, 100%	86.85%, 97.00%, 100%	91.57%, 98.22%, 100%	92.14%, 98.35%, 100%
Gumbel	85.15%, 100%, 100%	32.56%, 60.41%, 82.50%	35.35%, 54.51%, 76.20%	37.84%, 52.64%, 73.59%

"Table 2 about here"

Table 2: The Ratio of Seller's Revenue under ARMC to Seller's Revenue under Full Information

	[$\underline{v}_b$, $\bar{v}_b$] = [\$500K, \$1500K] v_s =\$750K, \$1000K, \$1500K	[$\underline{v}_b$, $\bar{v}_b$] = [\$1000K, \$2000K] v_s =\$750K, \$1000K, \$1500K	[$\underline{v}_b$, $\bar{v}_b$] = [\$1000K, \$2500K] v_s =\$750K, \$1000K, \$1500K	[$\underline{v}_b$, $\bar{v}_b$] = [\$1000K, \$3000K] v_s =\$750K, \$1000K, \$1500K
Exponential	80.46%, 100%, 100%	80.47%, 80.46%, 100%	74.89%, 74.39%, 100%	68.34%, 68.52%, 98.40%
Normal	97.87%, 100%, 100%	98.35%, 97.00%, 100%	98.92%, 98.22%, 100%	98.86%, 98.35%, 99.28%
Gumbel	98.24%, 100%, 100%	70.49%, 60.41%, 100%	64.49%, 54.51%, 100%	61.21%, 52.64%, 98.19%

As the figures in the Tables 1 and 2 suggest, the uniform distribution assumption performs well when the underlying distribution is normal. It may perform poorly for the exponential and Gumbel distributions, especially under very low capacity and moderate seller values. This is mainly because, if the underlying distribution is too skewed, the uniform distribution assumption yields a significant miscalculation in the value of the optimal bid. If the capacity is sufficiently large, the initial mishap could be remedied quickly as the bid given according to the uniform distribution assumption converges fast to the real optimal bid value, hence resulting in low revenue loss. If the seller valuation is too large, again the two revenue figures are close to each other, which is because buyers whose bids are accepted are almost the same regardless of the underlying distribution.

5.3 Stochastic Dynamic BPP Problem

Example 1 (continued): Recall the problem of a seller who has C=375 units to sell over T=15 time periods, where the buyers arrive with rate Λ=100 per period. The buyer valuation distribution is uniform on the range $[\mu(t) - \$300K, \mu(t) + \$300K]$ where $\mu(t) = \mu(t-1) + \delta(t)$, $\delta(t) = \{(-d, 0, d) \,\text{w.p.}\, (0.4, 0.2, 0.4)\}$ and $d \sim \mathbf{U}[0, \$120K]$. Suppose that $\mu(1)$=\$600K and the buyers bid according to the function $b^*(v_b) = 0.5 v_b$. In the base case, assume we do not account for the salvage value of the seller, hence $v_s = 0$.

We solve the scenario-based optimization problem for various seller valuation (v_s), capacity (C), and noise-size ($d := |\delta|$) values with N=150 scenarios; and compare the results with the simple "expected value (EV)" heuristic, where all stochastic variables in the problem are assumed to take their expected values, and with the solution of the uncertain QC-formulation given in (23)-(27), which we call "uncertain quadratically constrained (UQC)" solution. For all policies, we apply the proposed bid values on a random sample of 1000 scenarios, compute the revenues in all cases, and state the average of these revenues as a percentage of the absolute upper bound, i.e. the revenue produced by the "clairvoyant" policy. For the closed-loop policy we also compute the worst-case regret (z^*) within the scenarios used in the optimization model and state its ratio to the average revenues. We also investigate the effect of reformulating and resolving the problem at the beginning of each period according to current capacity and buyer valuations. All problems are solved via the CVX package developed by Grant and Boyd [17] for MATLAB using a version 7.5.0 and on a computer that has 4 GB of RAM. The results are given in the Tables 3, 4, 5 and 6[3].

"Table 3 about here"

Table 3: Changing v_s

	$v_s = 0$		$v_s = \$120K$		$v_s = \$180K$		$v_s = \$240K$	
	z^*	avg. r.	z^*	avg. r.	z^*	avg. r.	z^*	avg. r.
closed-loop	\$100K (6.2%)	93.28%	\$100K (8.6%)	91.21%	\$100K (10.7%)	89.32%	\$100K (14.2%)	86.24%
EV heuristic	-	70.81%	-	60.44%	-	57.72%	-	60.18%
UQC solution	-	43.12%	-	43.17%	-	44.17%	-	47.13%
closed-loop (res.)	-	96.36%	-	95.35%	-	93.98%	-	91.10%
EV (resolved)	-	90.17%	-	88.84%	-	87.62%	-	85.73%

"Table 4 about here"

Table 4: Changing C

	$C = 185$		$C = 375$		$C = 560$		$C = 750$	
	z^*	avg. rev	z^*	avg. rev	z^*	avg. rev.	z^*	avg. rev
closed-loop	\$77.5K (9.3%)	90.80%	\$100K (6.2%)	93.28%	\$107.2K (4.7%)	94.94%	\$108.7K (3.8%)	95.74%
EV heuristic	-	66.64%	-	70.81%	-	74.47%	-	78.51%
UQC solution	-	82.17%	-	43.12%	-	29.48%	-	22.80%
closed-loop (res.)	-	93.58%	-	96.36%	-	97.72%	-	97.97%
EV (resolved)	-	82.80%	-	90.17%	-	93.16%	-	94.85%

"Table 5 about here"

[3]Note that since the problem is resolved in each period, the z^* value which indicates the in-sample performance is not applicable for the iterative closed-loop policy.

Table 5: Changing d

	$d = \$60K$		$d = \$120K$		$d = \$240K$		$d = \$480K$	
	z^*	avg. r	z^*	avg. r	z^*	avg. r.	z^*	avg. r
closed-loop	\$48.6K (1.9%)	96.63%	\$100K (6.2%)	93.28%	\$384.5K (13.2%)	86.50%	\$894K (25.0%)	77.91%
EV heuristic	-	72.19%	-	70.81%	-	66.76%	-	59.93%
UQC solution	-	15.12%	-	43.12%	-	50.71%	-	50.68%
closed-loop (res.)	-	98.86%	-	96.36%	-	89.51%	-	80.93%
EV (resolved)	-	92.36%	-	90.17%	-	85.64%	-	78.92%

"Table 6 about here"

Table 6: Changing N

	$N = 50$		$N = 150$		$N = 200$		$N = 250$	
	z^*	avg. rev	z^*	avg. rev	z^*	avg. rev.	z^*	avg. rev
closed-loop	\$84K (5.9%)	92.78%	\$100K (6.2%)	93.28%	\$122K (7.6%)	92.85%	\$124K (7.7%)	93.07%

Here are a few remarks to note:

1. The worst case regret value is relatively low with respect to average revenues in almost all cases. This is an indication that the closed-loop formulation is not too conservative.

2. The closed-loop policy always outperforms the open-loop formulation of the uncertain QC problem, and the simple expected value heuristic. Moreover, there exist significant gains in resolving the problem at the beginning of each period with the current data (and possibly with more accurate future forecast figures).

3. The expected value heuristic also performs well if it is resolved at each period. This can be explained by the fact that feedback-type policies perform well if tracked in a smart manner; and is also in accordance with the findings of the literature, e.g. see Besbes and Maglaras [6] for a similar argument again regarding the real-estate sector.

4. The gap between the simple expected value heuristic (resolved) and the closed-loop heuristic (resolved) tends to be larger in the capacity-constrained settings. This is quite intuitive, since the scenario-based approach can account for various states of the world and prevent shortages; whereas the myopic approach does not have a pre-emptive nature.

5. Finally, both the in-sample (i.e. maximum regret) and out-of-sample (i.e. the revenue gap in other scenarios) performances of the closed-loop policy do not vary much by increasing N

after a critical number is reached. Moreover, this critical number of scenarios is expected to be as low as $N = 150$ for a problem of the above size.

6 Conclusion

In this paper, we discussed the dynamic negotiation problems, particularly in a buyer's market. We started with the one-to-one negotiation problems and discussed how to account for uncertainty in valuation distributions. Next, we extended our analysis to the dynamic environment: Starting with the deterministic fluid problem, we observed the stationary nature of optimal pricing policy. We were then able to extend the analysis to uncertain environments, and offer tractable and effective solution methodologies for real life applications.

Our results offer various avenues for future research: First, several other dynamic negotiation problems may be analyzed from the perspective we presented. Of these, the games that involve strategic buyers is of utmost interest. Also, the closed-loop formulation and the structural results regarding the nature of the optimal pricing policies might be inspiring and insightful in the formulation and solution of various other scenario-based robust optimization problems.

7 Appendix

7.1 Proof of Theorem 1

First, note that any optimal strategy should satisfy $b(v_b) \leq v_b$ and $s(v_s) \geq v_s$ to be feasible. This, combined with the assumption that the optimal strategies are nondecreasing in the valuations of the bidders, will be our implicit assumptions throughout the analysis and will be shown to hold.

In the minimax absolute regret minimization problem (7) of the seller, the innermost maximization takes the following values depending on the relationship among b, s and v_s:

$$\max_{s'}[(kb + (1-k)s' - v_s)1_{\{b \geq s'\}} - (kb + (1-k)s - v_s)1_{\{b \geq s\}}]$$
$$= \begin{cases} 0 & \text{if } b < v_s \\ (b - v_s) & \text{if } v_s \leq b \leq s \\ (b - v_s) - (kb + (1-k)s - v_s) & \text{if } b > s \end{cases}$$

That is, if the buyer bid is less than the seller's valuation, then any feasible bid of the seller returns zero net profit. If, the buyer bid exceeds v_s, the seller achieves her maximum profit by selecting the same bid as the buyer; which is the situation in the second and third cases in the above equivalence. Observe that in the second case, the seller overbids; whereas in the last case, she underbids and loses additional revenue she could have obtained if she had increased her bid up to b. Adding the outside maximization problem, the mathematical quantity to be minimized by selecting s is:

$$\begin{aligned} &\max_{b} \max_{s'}[(kb + (1-k)s' - v_s) \cdot 1_{\{b \geq s'\}} - (kb + (1-k)s - v_s) \cdot 1_{\{b \geq s\}}] \\ &= \begin{cases} 0 & \text{if } b < v_s \\ (s - v_s) & \text{if } v_s \leq b \leq s \\ (1-k)(\bar{b} - s) & \text{if } b > s \end{cases} \\ &= \max\{(s - v_s), (1-k)(\bar{b} - s)\} \end{aligned} \tag{32}$$

where $\bar{b}$ is the unknown maximum value of the buyer's bid b. Thus, the problem of the seller pours into selecting the bid to minimize the maximum of two regret values: In situation 1, the regret stems from overbidding and losing the chance to obtain positive return; whereas in situation 2, it stems from bidding too low and losing the chance of higher profits.

Since the first of the quantities inside the maximization in (32) is increasing and the second is decreasing in s, the minimizer is attained at the intersection point, i.e:

$$\begin{aligned} & s^*_{ARMC}(v_s) = \operatorname*{argmin}_{s} \max\{(s - v_s), (1-k)(\bar{b} - s)\} \\ \Rightarrow \quad & s^*_{ARMC}(v_s) - v_s = (1-k)(\bar{b} - s^*_{ARMC}(v_s)) \\ \Rightarrow \quad & s^*_{ARMC}(v_s) = \frac{v_s}{2-k} + \frac{(1-k)}{2-k}\bar{b} \end{aligned}$$

Via a symmetrical analysis for the buyers, we obtain $b^*_{ARMC}(v_b) = \frac{v_b}{1+k} + \frac{k}{1+k}\underline{s}$. Finally, since s^* and b^* should be best responses to each other, we find that $s^*_{ARMC}(v_s) = \frac{v_s}{2-k} + \frac{(1-k)\bar{v}_b}{2} + \frac{k(1-k)\underline{v}_s}{2(2-k)}$ and $b^*_{ARMC}(v_b) = \frac{v_b}{1+k} + \frac{k\underline{v}_s}{2} + \frac{k(1-k)\bar{v}_b}{2(1+k)}$. Furthermore, when the equations (1) and (2) are solved simultaneously for a game where both F_s and F_b are uniform, the resulting equilibrium bidding functions are identical to $s^*_{ARMC}(v_s)$ and $b^*_{ARMC}(v_b)$. $\Box$

7.2 One-to-one Negotiation Problem between Informed and Uninformed Agents

We know the solution to the one-to-one negotiation problem when (i) both the seller and the buyer know each other's distribution function, (ii) neither the seller nor the buyer know each other's distribution function, but they know the support of this function and employ ARMC approach to decide their bid. In this note, we will analyze a third case: (iii) the seller knows the buyer distribution function, F_b, while the buyer only knows the seller value range, $[\underline{v}_s, \bar{v}_s]$ (or, vice-versa).

During the analysis, we will implicitly assume that $s(.)$ and $b(.)$ are increasing in the seller and the buyer valuations respectively. At the end, we will show that this claim is true, provided that F_b is a distribution function with decreasing hazard rate (DFR).

The revenue maximization problem of the seller takes the form:

$$\begin{aligned}
\mathit{\Pi}_s(s, v_s) &= \max_{s\in[v_s,\bar{b}]} \int_s^{\bar{b}} (kb + (1-k)s - v_s) g_b(b) db, \\
&= \max_{s\in[v_s,b(\bar{v}_b)]} \int_{b^{-1}(s)}^{\bar{v}_b} (kb(v_b) + (1-k)s - v_s) f_b(v_b) dv_b,
\end{aligned}$$

which is maximized at the value s that satisfies the following equation:

$$(1-k)(1-F_b(b^{-1}[s]))b'(b^{-1}[s]) - f_b(b^{-1}[s])(s - v_s) = 0 \tag{33}$$

At this point, the seller does not know the function $b(.)$, or its derivative $b'(.)$. Thus, we turn our attention to the buyer's problem, which takes the form:

$$\operatorname*{argmin}_b \Big\{ \max_s \max_{b'} \left[(v_b - (kb' + (1-k)s)) \cdot 1_{\{b' \geq s\}} - (v_b - (kb + (1-k)s)) \cdot 1_{\{b \geq s\}}\right] \Big\} \tag{34}$$

$$= \operatorname*{argmin}_b \Big\{ \max\{(v_b - b), k(b - \underline{s})\} \Big\} \tag{35}$$

given that the buyer employs ARMC approach, by the analysis in the proof of Theorem 1.

Observe that, if the buyer is able to characterize the value of the lowest seller bid, $\underline{s} = s(\underline{v}_s)$, the solution of the equation (35) leads to the following bidding function:

$$b(v_b) = \frac{v_b}{k+1} + \frac{k\underline{s}}{k+1} \tag{36}$$

Hence, the seller's problem is equivalent to finding the s value that satisfies:

$$(1-k)(1-F_b(s(k+1)-k\underline{s}))\frac{1}{k+1} - f_b(s(k+1)-k\underline{s})(s-v_s) = 0$$

by inserting the appropriate values of b and b' into the equation (33).

Finally, note that the value of $\underline{s}$ is found from the equation:

$$(1-k)(1-F_b(\underline{s}))\frac{1}{k+1} - f_b(\underline{s})(\underline{s}-\underline{v}_s) = 0$$

which is then used to characterize the final form of the function $b(.)$. This final part is only true if $s(.)$ is nondecreasing in v_s, and a sufficient (but not necessary) condition to ensure this is that F_b is a function with decreasing failure rate. A symmetrical problem can be solved for the case where the buyer knows F_s while the seller only knows $[\underline{v}_b, \bar{v}_b]$. $\square$

7.3 Proof of Theorem 3

As before, our implicit assumptions are that the optimal strategies satisfy $b(v_b) \leq v_b$ and $s(v_s) \geq v_s$; and that the optimal strategies are nondecreasing in the valuations of the bidders.

Since buyers are naive, their problem takes the form:

$$\underset{b}{\operatorname{argmin}}\{\max_{s}\max_{b'}[(v_b-(kb'+(1-k)s))\cdot 1_{\{b'\geq s\}} - (v_b-(kb+(1-k)s))\cdot 1_{\{b\geq s\}}]\} \tag{37}$$
$$= \max\{(v_b-b), k(b-\underline{s})\} \tag{38}$$

As they assume that the seller is playing a one-to-one game with them, they simply compute their optimal bidding strategy by solving the two ARMC problems simultaneously, therefore reaching at the equilibrium bidding function of the one-to-one game, i.e. b^*_{ARMC}.

However, the seller's problem is now different: Given that buyers bid according to b^*_{ARMC}, she should select bid $s_t = s$, $\forall t$, that minimizes her maximum regret for all distribution functions F_b:

$$\begin{aligned}\underset{s}{\operatorname{argmin}}\{\max_{F_b}\max_{s'} \quad & \Big[\int_{t=0}^{T} \Lambda_t[\int_{b^{-1}(s')}^{\bar{v}_b}(kb(v_b)+(1-k)s'-v_s)f_b(v_b)dv_b]dt \\ & -\int_{t=0}^{\min\{T,T'\}} \Lambda_t[\int_{b^{-1}(s)}^{\bar{v}_b}(kb(v_b)+(1-k)s-v_s)f_b(v_b)dv_b]dt\Big]\}\end{aligned}$$

where s' is such that $\int_0^T \Lambda_t \int_{b^{-1}(s')}^{\bar{v}_b} f_b(v_b)dv_b dt = C$; and T' is such that $\int_0^{T'} \Lambda_t \int_{b^{-1}(s)}^{\bar{v}_b}[f_b(v_b)dv_b]dt = C$, if $s < s'$.

Regarding the inner maximization problem, we have two cases:

Case (i): $s < s'$**:** In this case the seller underbids and fails to capture a higher profit. The loss is at its maximum when all buyers have the highest valuation, i.e. $f_b(\bar{v}_b) = 1$. Thus:

$$\begin{aligned}&\max_{F_b}\max_{s'}\Big\{\int_{t=0}^{T} \Lambda_t[\int_{b^{-1}(s')}^{\bar{v}_b}(kb(v_b)+(1-k)s'-v_s)f_b(v_b)dv_b]dt \\ &-\int_{t=0}^{\min\{T,T'\}} \Lambda_t[\int_{b^{-1}(s)}^{\bar{v}_b}(kb(v_b)+(1-k)s-v_s)f_b(v_b)dv_b]dt\Big\} \\ &=\int_{t=0}^{T}\Lambda_t[(kb(\bar{v}_b)+(1-k)(b(\bar{v}_b))-v_s)]dt-\int_{t=0}^{T}\Lambda_t[(kb(\bar{v}_b)+(1-k)s-v_s)]dt \\ &=((1-k)(b(\bar{v}_b)-s))\min\{C,\int_{t=0}^{T}\Lambda_t\}dt\end{aligned}$$

Case (ii): $s > s'$**:** In this case the seller overbids and fails to sell a proportion of her inventories. This loss is at its maximum when all buyers bid just slightly below the seller's bid s, i.e. $f_b(b^{-1}(s - \epsilon)) = 1$ for small $\epsilon > 0$. Thus, the two inner maximization problems take the form:

$$\begin{aligned}&\max_{F_b}\max_{s'}\Big\{\int_{t=0}^{T} \Lambda_t[\int_{b^{-1}(s')}^{\bar{v}_b}(kb(v_b)+(1-k)s'-v_s)f_b(v_b)dv_b]dt \\ &-\int_{t=0}^{T} \Lambda_t[\int_{b^{-1}(s)}^{\bar{v}_b}(kb(v_b)+(1-k)s-v_s)f_b(v_b)dv_b]dt\Big\} \\ &=\int_{t=0}^{T}\Lambda_t[(k(s-\epsilon)+(1-k)(s-\epsilon)-v_s)]dt-0=(s-v_s)\min\{C,\int_{t=0}^{T}\Lambda_t\}\end{aligned}$$

Combining the two cases, the seller should bid to minimize the two maximum regrets, i.e. $s =$

argmin max$\{(s-v_s)\min\{C,\int_{t=0}^{T}\Lambda_t\},(1-k)(\bar{b}-s)\min\{C,\int_{t=0}^{T}\Lambda_t\}\}$. But these two regret terms are the same terms as in the one-to-one game, only multiplied by a coefficient $\min\{C,\int_{t=0}^{T}\Lambda_t\}$. Thus, we arrive at the same conclusion as before; i.e. the seller bids as if F_b is uniform on its given range, which also validates the buyers' bidding game. □

References

[1] A. Nemirovski A. Ben-Tal and C. Roos. Robust solutions of uncertain quadratic and conic-quadratic problems. *SIAM Journal on Optimization*, 13(2):535–560, 2002.

[2] E. Guslitzer A. Ben-Tal, A. Goryashko and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351376, 2004.

[3] S. Boyd A. Ben-Tal and A. Nemirovski. Control of uncertainty-affected discrete time linear systems via convex programming, 2005.

[4] M. O. Ball and M. Queyranne. Toward robust revenue management: Competitive analysis of online booking. *Operations Research*, 57(4):950–963, 2009.

[5] D. Bertsimas and D. B. Brown. Constrained stochastic lqc: A tractable approach. *IEEE Transactions on Automatic Control*, 52(10):18261841, October 2007.

[6] O. Besbes and C. Maglaras. Dynamic pricing with financial milestones: Feedback-form policies, 2009. Working Paper.

[7] A. Bhandari and N. Secomandi. Revenue management with bargaining. *Operations Research*, 59(2):408–506, 2011.

[8] K. Chatterjee and W. Samuelson. Bargaining under incomplete information. *Operations Research*, 31(5):835–861, September-October 1983.

[9] S. Clifford. The ny times article in the adress http://www.nytimes.com/2012/03/28/business/retailers-rush-to-adjust-to-price-smart-shoppers.html, 2012.

[10] D. A. Iancu D. Bertsimas and P. A. Parrilo. Optimality of affine policies in multi-stage robust optimization. *Mathematics of Operations Research*, 35:363–394, 2010.

[11] S. Eren and C. Maglaras. Monopoly pricing with limited demand information. *Journal of revenue and pricing management*, 9(1-2):23–48, 2010.

[12] S. Eren and G. van Ryzin. Product line positioning without marked information. *Working Paper, Columbia University, New York, NY 10027*, 2006.

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

[14] J. Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54:291–310, 2006.

[15] Stanley J. Garstka and Roger J.-B.Wets. On decision rules in stochastic programming. *Mathematical Programming*, 7(1):117–143, 1974.

[16] D. Goldfarb and G. Iyengar. Robust convex quadratically constrained programs. *Mathematical Programming*, 97(3):495–515, 2003.

[17] M. Grant and S. Boyd. *CVX: MATLAB Software for Disciplined Convex programming*. 2008.

[18] J. C. Harsanyi. Approaches to the bargaining problem before and after the theory of games. *Econometrica*, 24(2):144–157, April 1956.

[19] E.C. Kerrigan and J.M. Maciejowski. Properties of a new parameterization for the control of constrained systems with disturbances. *Proceedings of the 2004 American Control Conference*, 5(5):46694674, June-July 2004.

[20] Lan, Gao, M. Ball, and I Karaesmen. Revenue management with limited demand information. *Working Paper, University of Maryland, College Park, MD*, 2006.

[21] T. C. Lee and M. Hersh. A model for dynamicairline seat inventory control with multiple seat bookings. *Transportation Science*, 27(3):252–265, august 1998.

[22] Lobel and G. Perakis. Network dynamic pricing through scenario based optimization, 2010. Working Paper.

[23] C. Maglaras and J. Meissner. Dynamic pricing strategies for multi-product revenue management problems. 8(2):136–148, 2006.

[24] R. B. Myerson. Two person bargaining problems with incomplete information. *Econometrica*, 2(2):461–487, March 1984.

[25] R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

[26] J. C. Nash. The bargaining problem. *Econometrica*, 18:151–162, 1950.

[27] G. Perakis and G. Roels. Robust controls for network revenue management. *Working Paper, UCLA, Los Angeles, CA 90095*, 2007.

[28] J. Riley and R. Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *Quarterly Journal of Economics*, 98(2):267–289, 1983.

[29] J. Sinclair. The internet blog in the adress http://jsmineset.com/, 2009.

[30] G. Das Varma and N. Vettas. Optimal dynamic pricing with inventories. *Economic Letters*, 72:335–340, 2001.

[31] G. Vulcano, G. van Ryzin, and C. Maglaras. Optimal dynamic auctions for revenue management. *Management Science*, 45(11):1388–1407, November 2002.