Omega ■ (■■■■) ■■■–■■■

Contents lists available at ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

Revenue management with minimax regret negotiations ☆

Nur Ayvaz-Cavdaroglu [a,*], Soulaymane Kachani [b], Costis Maglaras [c]

[a] Kadir Has University, Department of Business Administration, Istanbul 34083, Turkey
[b] Columbia University, IE&OR Department, School of Engineering and Applied Sciences Columbia University, 500 West 120th Street, Rm. 313, New York, NY 10027, United States
[c] Columbia Business School, 409 Uris Hall, 3022 Broadway, New York, NY 10027, United States

ARTICLE INFO

Article history:
Received 2 February 2014
Accepted 25 September 2015

Keywords:
Allocation
Bidding
Resource management

ABSTRACT

We study the dynamic bilateral price negotiations from the perspective of a monopolist seller. We first study the classical static problem with an added uncertainty feature. Next, we review the dynamic negotiation problem, and propose a simple deterministic "fluid" analog. The main emphasis of the paper is in analyzing the relationship of the dynamic negotiation problem and the classical revenue management problems; and expanding the 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 acceptable price at any given point in time.

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 [5,16,18]. There are also examples in consumer commerce [34,19,15,7,10,30]. The outcome of each such negotiation depends on the reservation values of the seller and buyer, their negotiation skills, and their beliefs on the same parameters of their opponent. This process is known as a "bilateral negotiation", and if the focus of the negotiation process is restricted to prices specifically, as "bilateral price negotiations".

Despite the importance and prevalence of negotiation problems in practice, quantitative dynamic pricing and revenue management, which has "evolved into a mature research area to support a seller's tactical capacity allocation choices and pricing decisions with inventory considerations [24]" has mostly focused on posted price mechanisms [11,35] and auctions [36]. There have been several extensions of the classical revenue management problem, for instance Bodily and Weatherford [4] consider the situations with continuous resources and several pricing classes; Sen [32] develops dynamic pricing heuristics as an extension to the Gallego and Van Ryzin's model that perform substantiall\y better than the fixed price policy. Lan et al. [20,21] provide successful examples of combining the overbooking and seat allocation decisions with the regret models. (Among other interesting line of research lie Kim and Bell's work [17] on the optimal pricing and production decisions in the presence of substitution, Tsai and Hung's paper [33] on the use of integrated real options internet retailing, Zhao et al. [37] regarding dynamic pricing in the presence of customer inertia, and Ghoniem and Maddah [13] optimizing retail assortment, pricing, and inventory decisions with substitutable products.) However, this broad research area has largely ignored the bilateral price negotiation problems perhaps regarding them as being in the scope of game theory. However, as we emphasize in this paper, the two problem types could be very similar and revenue management methods can be readily applicable in bilateral negotiation problems.

In more detail, we hereby focus on 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. These buyers 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 engage 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 bilateral negotiation is modeled as a one-off negotiation, where the buyer and seller submit bids and where the unit is awarded if the buyer's bid is higher than the seller's bid. When the seller has market power, the transaction price is the seller's posted price (SPP); when the buyer

☆This manuscript was processed by Associate Editor M. Shen.
* Corresponding author.
E-mail address: nur.cavdaroglu@khas.edu.tr (N. Ayvaz-Cavdaroglu).

http://dx.doi.org/10.1016/j.omega.2015.09.009
0305-0483/

has market power, the transaction price is the buyer's posted price (BPP); in other cases the transaction price splits the difference between the two bids according to a fixed ratio that models the relative negotiation power of the two players.[1]

Among the papers that involve revenue management problems in the form of bilateral negotiations, the work of Bhandari and Secomandi [3] is perhaps closest to ours regarding the problem under consideration. However, the authors use a stylized MDP to investigate the negotiation processes and measure the performances of the seller under various negotiation mechanisms via numerical studies, while we resort to fluid approximation and develop an analytical result. Still, our findings in the numerical analysis section has common elements with their work. 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 [31] and Gallien [12]. Furthermore, Huang et al. [15] and Chen et al. [7] study the two selling mechanisms, namely, "posted price" versus "name-your-own-price" in a retail environment; however, the existence of several competing sellers, the forward-looking customers and other details differentiate their models from the model of our paper.

Finally, Kuo et al. [19] study a very similar problem to ours in the sense that they focus on retailers for whom take-it-or-leave-it price is the main mode of operation, but who nonetheless allow price negotiation when they encounter "bargainers". The retailer, as in the dynamic setting of our paper, encounters a series of bargainers over time, and the outcome of the negotiation with each bargainer will depend on the retailer's inventory and the remaining time until the end of the selling season. Their formulation differs from ours in how the outcome of each negotiation is characterized: in their work, the retailer sets a posted price, which acts as a ceiling on the revenue obtained from buyers, and a cutoff price which affects the final price according to the general Nash bargaining solution (GNBS); while we adopt the Chatterjee and Samuelson's model in which the seller sets a single bid value. Therefore, each party's lack of information about each other's valuation does not create a problem in their setting in terms of reaching a bargaining outcome, while the assumption in the classical one-to-one negotiation problem we consider is both parties having perfect information about each other's valuation distribution, which we happen to relax in the course of the paper. Moreover, the main focus of Kuo et al. is to study the effects of negotiation on the retailer's dynamic prices and revenues and the payments of both bargainers and price-takers in a variety of settings; while the ultimate focus of our paper is to study the classical and the dynamic bilateral negotiation problem with various extensions and to create a link between the economics and revenue management literatures by establishing its connection with the classical revenue management problems.

The first modeling and methodological contribution of our 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 [1], Eren and Maglaras [8], Perakis and Roels [29], Lan et al. [22] and Eren and Van Ryzin [9] adopt different versions of this idea. Adopting the maximum regret criterion, we formulate jointly the buyer and seller bidding problems in the setting where the underlying distribution functions F_s, F_b are unknown to the respective counterparties, and show that the optimal strategies are to bid as if these distributions were uniform. This result, to our knowledge, is novel in the literature; although there are several papers that accept that the players will *de facto* believe that the other party has uniform distribution and act on it.

Secondly, we turn our attention 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 [23]. 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 [11].

Next, motivated by the goal of studying the dynamic settings, 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. Furthermore, extending the findings of the one-to-one problem with the added uncertainty feature, it becomes possible to study a setting where the distributional assumptions are not known.

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 analog 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. The numerical analysis section complements the analytical findings from other interesting perspectives, namely investigating the effect of the negotiation power and the effect of the uniform distribution assumption on the revenues of the seller and the bids of the two parties.

1.1. The remainder of the paper

In Section 2, we analyze a variant of the classical one-to-one negotiation problem with an added uncertainty element. 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

[1] A detailed definition of each negotiation mechanism can be found in Bhandari and Secomandi [3].

Please cite this article as: Ayvaz-Cavdaroglu N, et al. Revenue management with minimax regret negotiations. Omega (2015), http://dx.doi.org/10.1016/j.omega.2015.09.009

a dynamic setting using a fluid model approach. Next, in Section 3.3, the results of Section 2 are extended to the dynamic setting under a regret criterion. Numerical illustrations and extensions are presented in Section 4. Finally, Section 5 concludes our findings and presents avenues for further research.

2. An extension to the 1-to-1 bilateral negotiation problem

The literature of two-person bargaining games goes back to Nash [28] and Harsanyi [14], and the ones to pioneer the analysis of the dynamics of an environment with shifting negotiation power are Myerson (et al.) [25–27] and Chatterjee and Samuelson [6]. 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 $[\underline{v_b},\overline{v_b}]$. 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 $[\underline{v_s},\overline{v_s}]$. F_s and F_b are both strictly increasing and differentiable on their supports, and are common knowledge.

The rules of the bargaining game are 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 $s(v_s) \le b(v_b)$. 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 $s(v_s) \le b(v_b)$ holds. In general, the equilibrium of the game is found by solving the following "best response problems" of the seller and the buyer simultaneously:

$$\max_{s \epsilon \left[v_s,\overline{b}\right]} \int_s^{\overline{b}} (kb+(1-k)s-v_s)g_b(b)db,$$

and

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

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 $\overline{b}$ is the maximum value the buyer's bid can assume.

Chatterjee and Samuelson [6] characterize the class of equilibria for the above game in which player bidding strategies are "well-behaved". In particular, they make the following assumption regarding the buyer and the seller bidding functions $s(.)$ and $b(.)$, which is also relevant for our analyzes:

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] happen to be the solutions to the following two linked differential equations:

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

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

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

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

$$b^*(v_b)=\{b| -F_s(b)+(v_b-b)f_s(b)=0\} \quad \forall v_b \in \left[\underline{v_b},\overline{v_b}\right], \quad (3)$$

$$s^*(v_s)=v_s, \quad \forall v_s \in \left[\underline{v_s},\overline{v_s}\right], \quad (4)$$

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

$$b^*(v_b)=v_b, \quad \forall v_b \in \left[\underline{v_b},\overline{v_b}\right] \quad (5)$$

$$s^*(v_s)=\{s| 1-F_b(s)+(v_s-s)f_b(s)=0\}, \quad \forall v_s \in \left[\underline{v_s},\overline{v_s}\right] \quad (6)$$

As evident from these sets of equations, the seller's optimal bidding function is independent from the buyer's value distribution function in the BPP setting and vice versa in the SPP setting. The intuition behind this fact is obvious: under the BPP mechanism the seller has no influence on determining the final price, therefore 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 becomes 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).

The classical one-to-one bilateral negotiation problem is famous and explicitly analyzed in the literature. However, the following variant of the problem with an added uncertainty feature is, to our knowledge, not specifically discussed. 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 respectively as follows:

$$\underset{s}{\text{argmin}}\left\{\max_b \max_{s^0}\left[(kb+(1-k)s^0-v_s).1_{\{b \ge s^0\}}-(kb+(1-k)s-v_s).1_{\{b \ge s\}}\right]\right\}=0 \quad (7)$$

$$\underset{b}{\text{argmin}}\left\{\max_s \max_{b^0}\left[\left(v_b-kb^0-(1-k)s\right).1_{\{b^0 \ge s\}}-(v_b-kb-(1-k)s).1_{\{b \ge s\}}\right]\right\}=0 \quad (8)$$

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^0 (i.e. $(kb+(1-k)s^0-v_s)\cdot 1_{\{b \ge s^0\}}$) and the realized revenue at her selected bid s (i.e. $(kb+(1-k)s-v_s)\cdot 1_{\{b \ge s\}}$). 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.

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

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:*

$$s^*_{ARMC}(v_s) = \frac{v_s}{2-k} + \frac{(1-k)\overline{v_b}}{2} + \frac{k(1-k)\underline{v_s}}{2(2-k)}, \quad \forall v_s \in \left[\underline{v_s}, \overline{v_s}\right] \tag{9}$$

$$b^*_{ARMC}(v_b) = \frac{v_b}{1+k} + \frac{k\underline{v_s}}{2} + \frac{k(1-k)\overline{v_b}}{2(1+k)}, \quad \forall v_b \in \left[\underline{v_b}, \overline{v_b}\right] \tag{10}$$

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 A.1 Appendix A.

The above result brings a theoretical motivation to use uniform distribution as the opponent's distribution function when there is no relevant 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. This is a very widely used approximation in the literature, since Myerson [26] introduced the famous SUTP (symmetric uniform trading problem) where both parties believe that their opponent's distribution is uniform on the range [0,1]. However, other than regarding it as a simple de facto belief, the theoretical roots of the uniform distribution assumption had not been analyzed in the literature before. Next, we extend this result to a dynamic setting and complete the analysis in the numerical results section by exploring the effect of this assumption on the seller revenues.

3. Dynamic 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 (which might or might not be changing with time), 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.

This problem is extensively analyzed in the revenue management literature, but not taking into account the negotiation perspective. In the following subsection, we establish the connection of the dynamic negotiation problem with the classical problems of the revenue management literature. Next, we develop the solution of the problem by using fluid formulation and extend it to an uncertain environment by combining the elements of the one-to-one setting and the fluid formulation.

3.1. The analogy of revenue management and bilateral negotiation problems

In this part, we focus on the connection between the dynamic bilateral negotiation problem and the revenue management problems.

First, consider the SPP (seller posted price) setting: in this setting, given the arrival rate Λ_t and the buyer bidding function $b^*_{SPP}(v_b) = v_b$, the problem of the seller is to find the best dynamic pricing policy that would maximize the total net revenues of the seller, which is the cumulative of the net revenues extracted from each successful individual negotiation. Given the seller's bid s_t, the demand process can be regarded as a one-dimensional nonhomogeneous Poisson process with rate vector λ determined through a demand function $\lambda(s_t)$ where $\lambda(s_t) = \Lambda_t(1-F_b(s_t))$, with a continuous inverse demand function $p(\lambda) = s_t$, since the buyers pay the seller's bid in SPP environment regardless of their value. Then, adopting a discrete-time formulation (which assumes that time has been discretized in small intervals of length δt indexed by $t=1,\ldots,T$ such that $P(\text{arrival in } [0,\ \delta t]) = \lambda\delta t + o(\delta t)$ and $P(\text{two arrivals in } [0,\ \delta t]) = (\lambda\delta t)^2 + o((\delta t)^2)$, where $o(x)$ implies that $o(x)/x \to 0$ as $x \to 0$); the corresponding demand random variable for period t, denoted by $\xi(t;\lambda)$, is Bernoulli with probabilities that are controlled by the posted price (i.e. seller's bid s_t), and $P(\xi(t)=1) = \lambda(s_t)\delta t$ and $P(\xi(t)=0) = 1-\lambda(s_t)\delta t$. In this environment, the seller's revenue maximization problem could be formulated as follows:

$$\max_{\{s_t,\ t=1,\ldots,T\}} E_\xi\left[\sum_{t=1}^{T}(s_t - v_s)\xi(t;\lambda)\right] \tag{11}$$

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

where P denotes the feasible price set. This formulation is mathematically equivalent to the famous "dynamic pricing problem" of a monopolist seller selling a homogenous product in a discrete-time setting; which is readily analyzed by Gallego and Van Ryzin [11] (which will be referred to as GVR paper in the sequel). For further details of this formulation, please see Maglaras and Meissner [24].

Similarly, consider the BPP (buyer posted price) setting: in this environment, given the buyer bidding function b^*_{BPP}, it is possible to define the expected net revenue from each prospective negotiation at instant t as $p(s_t) = b_t - v_s$ where $b_t := b_{BPP}{}^*(v_t)$, provided that $b_t \ge s_t$. That is, the seller bid s_t has no direct effect on the exogenously fixed product prices, and a priori fixed prices also fix the demand rate. The seller can only control which product requests (given that the products in our example refer to the same unit with different prices, this is equivalent to determining which buyer bids) to accept at any given time. 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 [23] 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 with n finite values; i.e. define $\overline{b^*} \ge b_1 \ge b_2 \ge \ldots \ge b_n \ge \underline{b^*}$ as n finite "fare classes", where the arrival rate of bid b_i is approximated by $\Lambda_t\left(\overline{G_b}(b_i) - \overline{G_b}(b_{i-1})\right)$, &for all;$i \in \{2,\ldots,n\}$ and that of b_1 is approximated by $\Lambda_t\overline{G_b}(b_1)$ at each instant t, where $\overline{G_b}(s_t)$ is the complement of the cumulative distribution function (cdf) of b^*_{BPP}. 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:

$$\max_{\{u(t),\ t=1,\ldots,T\}} E_\xi\left[\sum_{t=1}^{T}(b' - v_s)\xi(t;u\lambda)\right]$$

$$\textit{subject to} \quad \sum_{t=1}^{T} e'\xi(t;u\lambda) \le C \text{ a.s.}, u_i(t) \in \{0,1\}, \forall t \tag{13}$$

where $u_i(t)$'s are the controls that take a value of one when a bid of value $b_i$is accepted at time t and zero otherwise (i.e. $u_i(t)=1$ if $b_i \ge s_t$ and $u_i(t)=0$ otherwise), $b' = \{b_1, b_2, \ldots, b_n\}$, e' the n-dimensional unit vector, and $\xi(t;u\lambda)$ is the corresponding demand random variable for period t with coordinates $u_i\lambda_i$, where $P(\xi(t;u_i\lambda_i) = 1) = u_i(t)\delta t[\Lambda_t\left(\overline{G_b}(b_i) - \overline{G_b}(b_{i-1})\right)]$, $P(\xi(t;u_i\lambda_i)=0) = 1 - u_i(t)\delta t[\Lambda_t\left(\overline{G_b}\right.$

$(b_i) - \overline{G_b}(b_{i-1}))]$ &for all;$i \in \{2,\ldots,n\}$ and $P(\xi(t;u_1\lambda_1)=1)=u_1(t)\delta t[\Lambda_t\overline{G_b}(b_1)]$, $P(\xi(t;u_1\lambda_1)=0)=1-u_1(t)\delta t[\Lambda_t\overline{G_b}(b_1)]$, The formulation (13) is equivalent to the discretized version of the capacity control problem of Lee and Hersh [23]. 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 further details on the formulation, we refer the reader to Maglaras and Meissner [24].

Both of the above equivalences stem from the fact that each buyer is *naive* (i.e. does not strategize on the time of purchase, does not take into account previous sales prices, etc.). Moreover, while each buyer negotiates with the seller only once, the seller will engage in a sequence of negotiations over the sales horizon. Hence, in SPP setting, the seller will pursue a *dynamic pricing strategy* to maximize the revenues to be extracted from the stochastically arriving buyers, whereas in BPP, she will determine the minimum bid to be accepted at each instant to control *the amount of capacity to be sold*. 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

Even when the buyers are naive, it is difficult to analyze the stochastic dynamic pricing problem of the seller. 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, one would typically solve the recursions numerically or resort to some approximations such as approximate dynamic programming or simulation. For instance, the authors have adopted a fluid formulation in the GVR paper and developed asymptotically optimal policies. In addition, in negotiation problems, there is a continuous stream of buyer bids often with their range and frequency varying in time; therefore it is practically impossible to approximate the buyer bids as finite number of fare classes as in Lee and Hersh [23]. Even if this approximation is valid, the issue of “curse of dimensionality” prevails. Thus, the size of the problem renders the computation of the optimal bids s_t, &for all;t, almost impossible, which leads us to develop a fluid formulation equivalent of the dynamic negotiation problems and focus on the analysis in this fluid setting, 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:

$$\max_{\{s_t,\ \forall t\}} \left[\int_{t=0}^{T} r_t(v_s, s_t)dt\right] \tag{14}$$

$$\text{subject to} \quad \int_{t=0}^{T} \Lambda_t \left[\int_{s_t}^{\bar{b}} g_b(b)db\right] dt \le C \tag{15}$$

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:

$$r_t(v_s, s_t) = \int_{s_t}^{\bar{b}} \Lambda_t(kb+(1-k)s_t - v_s)g_b(b)db \tag{16}$$

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 $\alpha t(st) = \int_{s_t}^{b} g_b(b)db = G_b(s_t)$, 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 (14) and (15) in the price space is equivalent to:

$$\max_{\{\alpha_t,\ \forall t\}} \left[\int_{t=0}^{T} r_{t,a}(v_s, \alpha_t)dt\right] \tag{17}$$

$$\text{subject to} \quad \int_{t=0}^{T} \Lambda_t \alpha_t dt \le C \tag{18}$$

which is a formulation in the demand space.

Provided that $r_{t,a}(v_s, \alpha)$ is concave in α[3], the formulation (17) and (18) 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:*

$$s_t(v_s) = \max\left\{G_b^{-1}\left(1 - \frac{C}{\int_{t=0}^{T}\Lambda_t dt}\right), s^*(v_s)\right\}, \quad \forall t \in [0, T] \tag{20}$$

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 alreadyoted, 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

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

$$g_b(b') \ge 0, \quad \forall b \in \left[\underline{b}, \bar{b}\right] \tag{19}$$

is a sufficient condition to ensure that $r_{t,a}(v_s, \alpha)$ is concave in α for all $v_s \in \left[\underline{v_s}, \overline{v_s}\right]$. 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 (19) is satisfied.

Please cite this article as: Ayvaz-Cavdaroglu N, et al. Revenue management with minimax regret negotiations. Omega (2015), http://dx.doi.org/10.1016/j.omega.2015.09.009

negotiation game and employ the equilibrium bidding function $b^*(.)$ regardless of their arrival time.

To see how the seller behaves, note that the problem (17) and (18) is maximized at the maximizer of $r_{t,a}(.,v_s)$, which is $\alpha^* = \overline{G_b}(s^*(v_s))$, provided that it is feasible to admit this fraction at each instant t (i.e. if $\alpha^* \int_{t=0}^{T} \Lambda_t dt \le C$). This case is equivalent to applying the bid $s_t(v_s) = s^*(v_s)$, &for all;t.

If, on the other hand, $\alpha^* \int_{t=0}^{T} \Lambda_t dt > C$, then by the concavity of $r_{t,a}(.,v_s)$, it is optimal to admit the constant fraction $\propto_0 = \frac{C}{\int_{t=0}^{T} \Lambda_t dt}$ at each t. This second case corresponds to bidding $s_t(v_s) = G_b^{-1}\left(1 - \frac{C}{\int_{t=0}^{T} \Lambda_t dt}\right), \forall t \in [0,T]$. So the seller will set her reservation price as $s_t(v_s) = \max\left\{G_b^{-1}\left(1 - \frac{C}{\int_{t=0}^{T} \Lambda_t dt}\right), s^*(v_s)\right\}$.

The above theorem is in the same spirit as the Proposition 2 of GVR paper and forms the first major result of this section. Basically it states that the seller has a unique "acceptable bid", which remains unchanged whereas the final price of the good in each negotiation might vary depending on the bids of the buyers. The value of the seller's minimum acceptable bid depends on the bidding function of the buyers, namely the parameter k. Although we expect the minimum acceptable bid and the revenues of the seller to increase as k decreases (i.e. as we approach a SPP setting), this is not always the case. This interesting phenomenon is analyzed in further detail in the first part of the numerical analysis section.

3.3. Dynamic negotiation problems under uncertainty

In this part, we study a variant of the 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 only revealed at t. The players know each other's distribution range at each instant $t \in [0,T]$ (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, and we can state and prove the following Theorem, which emphasizes the analogy between the dynamic problem with no distribution information and the stochastic one-to-one problem.

Theorem 3. *The dynamic negotiation problem with infinitesimal buyers and unknown valuation distributions 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 B. The basic idea is that the seller bids as if the prospective buyers will have uniform distribution in the given range. In a fluid setting with an unvarying buyer valuation range, the result is a stationary bidding policy for the seller. Although this result is oversimplified, it is in the same spirit as the "feedback policy" of Besbes and Maglaras [2], who state that the optimal policy for a monopolist seller in a stochastic environment that is similar to our original problem can be written in feedback form that dynamically tracks the revenue/financial milestones. Similarly, the implication of the fluid formulation's optimal solution for a seller in the stochastic environment is to dynamically track the sales rate, and correct the bidding price as to keep the stationary sales rate of the fluid solution.

As careful readers should immediately notice, the above theorem relies on the assumption that there is no "learning effect" for the seller. That is, the seller cannot infer the value distribution function of the future buyers from the current distribution function or the instantaneous sales rate. Because if this were the case, the problem becomes trivial and once the seller infers F_b, she can employ the optimal pricing policy which is characterized in (20).

Although the assumption that "seller cannot learn from experience" 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. This situation should not come as idiosyncratic to the reader. There are several industries where the willingness-to-pay values of the buyers depend closely on macroeconomic terms, and especially in instable economies, these terms can fluctuate substantially from one time instance to another. In many real life situations the reality falls in-between: that is, to a certain degree, it is possible to infer the valuation distribution of future buyers by observing the current valuations. However, this is an avenue that could itself be a complete research project and for the sake of brevity, we leave the pursuit of this research direction to a future work.

4. Numerical results

4.1. The effect of the negotiation parameter

So far, we have characterized the general form of the seller bidding function regardless of the underlying negotiation setting. In this subsection, we would like to measure the effect of the

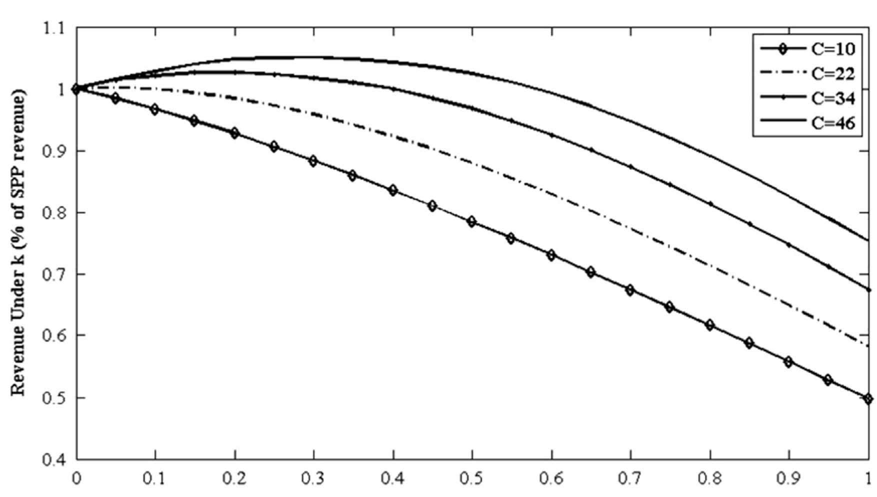

Fig. 1. Seller revenues (as % of revenue at $k=0$) for various k and C values.

Please cite this article as: Ayvaz-Cavdaroglu N, et al. Revenue management with minimax regret negotiations. Omega (2015), http://dx.doi.org/10.1016/j.omega.2015.09.009

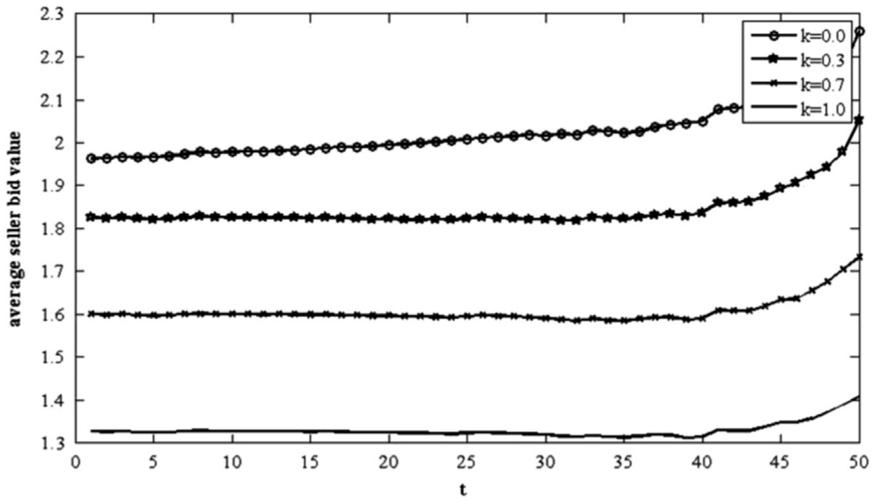

Fig. 2. Average seller bid values over time across 500 instances, for a capacity level of $C=26$, at different k values.

"buyer's negotiation power" (which is reflected in the parameter k) on the seller bids and revenues and on the prices paid by the buyers by a set of numerical experiments.

To this end, consider a dynamic setting where the buyers and the seller both have uniform valuation distributions on the ranges $[\underline{v_b}, \overline{v_b}]=[1,\ 3]$ and $[\underline{v_s}, \overline{v_s}]=[0.5, 1.5]$ respectively. Assume that the buyers arrive according to a Poisson distribution with rate $\Lambda=1$ per period through 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:

$$b^*(v_b)=\frac{v_b}{1+k}+\frac{k\underline{v_s}}{2}+\frac{k(1-k)\overline{v_b}}{2(1+k)}, \quad \forall v_b \in \left[\underline{v_b}, \overline{v_b}\right] \tag{21}$$

$$s_t^*(v_s)=\max\left\{\frac{v_s}{2-k}+\frac{(1-k)\overline{v_b}}{2}+\frac{k(1-k)\underline{v_s}}{2(2-k)}, G_b^{-1}\left(1-\frac{x(t)}{\int_{\tau=t}^{T}\Lambda_\tau d\tau}\right)\right\}$$
$$\forall t, \forall v_s \in \left[\underline{v_s}, \overline{v_s}\right] \tag{22}$$

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 in which the valuation of the seller, the number of arriving buyers, and the valuation of each arriving buyer are randomly determined according to the problem parameters. The ratio of average seller revenues for a given k value to the revenues under the SPP setting at various levels of seller capacity is depicted in the Fig. 1.

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 a high load factor might actually benefit from a slight shift in the negotiation power. This is an interesting observation, since the average seller bids tend to decrease in k and in the available inventory as expected (which can be seen in Figs. 2 and 3). However, considering the number of successful trades (as seen in Fig. 4), even a slight increase in k (e.g. from $k=0$ to $k=0.2$) is seen to produce a considerable increase in the average number of buyers that are accepted at high levels of capacity. This observation shows that the seller might actually benefit from "selling more at slightly reduced prices" which would be possible by shifting a small proportion of the negotiation power to the other party. This result is also consistent with the results of Bhandari and Secomandi [3], who find that "in structured negotiations a seller with an exceptionally high relative inventory availability can benefit from splitting the difference between his offer and that of each arriving buyer, rather than making a first and final offer", and, in a broader sense, with the findings of Kuo et al. [19] who claim that "the retailer's benefit from allowing negotiation increases in the retailer's bargaining power (…) which is particularly helpful when the inventory level is high relative to the remaining selling season". Moreover, although their bids are lower, the prices paid by the high-valued buyers (with valuations above 2.5) do not decrease much on average for low-to-moderate k values (e.g. $k=0.2$ or $k=0.3$) compared to the $k=0$ setting (Fig. 5). This result could be attributed to the fact that the seller's bid increases towards the end of the sales horizon because of a fast initial sale at low-to-moderate k values, and the final price paid by the late-arriving, high-valued buyers becomes similar to, or even higher than the prices paid under the SPP setting. This observation also strengthens our initial conclusion that the seller might benefit from selling more at slightly reduced prices initially, which is the result of sharing a small portion of the negotiation power with the other party.

Another interesting observation is the fact that 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 $k=0$; $b(v_b)=\frac{v_b}{1.2}+\frac{0.5x0.2}{2}+\frac{0.2x0.8x3}{2x1.2}=1.25$ for $k=0.2$, and $b(v_b)=\frac{v_b}{2}+\frac{0.5}{2}=0.85$ for $k=1$) This effect can be seen more explicitly in Fig. 6, which shows the bid values of buyers with various valuations. However, considering that the lower-valued buyers are less likely to conclude a successful trade with the seller (as seen in Fig. 5), the effect of this observation on the total seller revenues does not seem to be significant unless the seller has a much lower valuation for the good with respect to all buyers.

4.2. 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, we consider the revenue maximization problem of a seller who operates in a BPP setting, where the market size is Poisson with rate $\Lambda=100$ per period for $T=15$ periods.[4]

For the Normal and Gumbel distributions, we extracted the mean as the midpoint of the range and selected the standard

[4] Similar results were obtained for a problem in an SPP setting, but we do not consider it necessary to involve the results here for the sake of brevity.

Please cite this article as: Ayvaz-Cavdaroglu N, et al. Revenue management with minimax regret negotiations. Omega (2015), http://dx.doi.org/10.1016/j.omega.2015.09.009

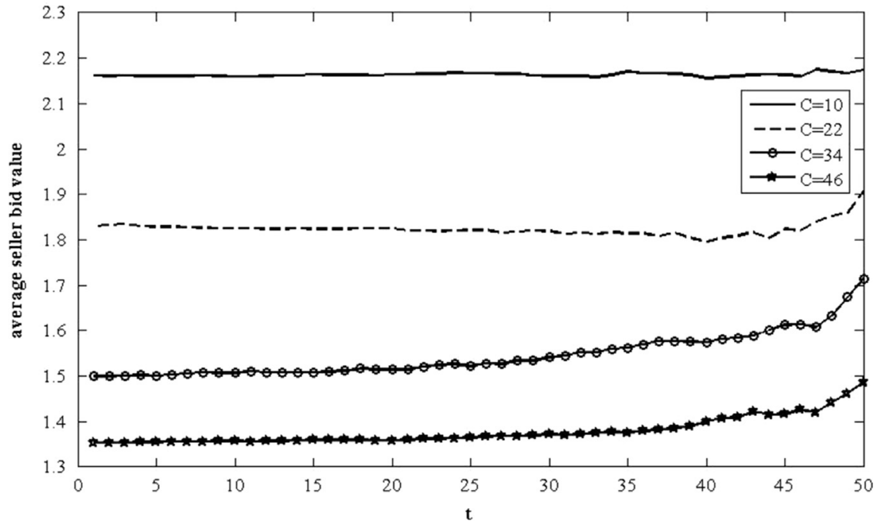

Fig. 3. Average seller bid values over time across 500 instances at $k=0.5$, at different C levels.

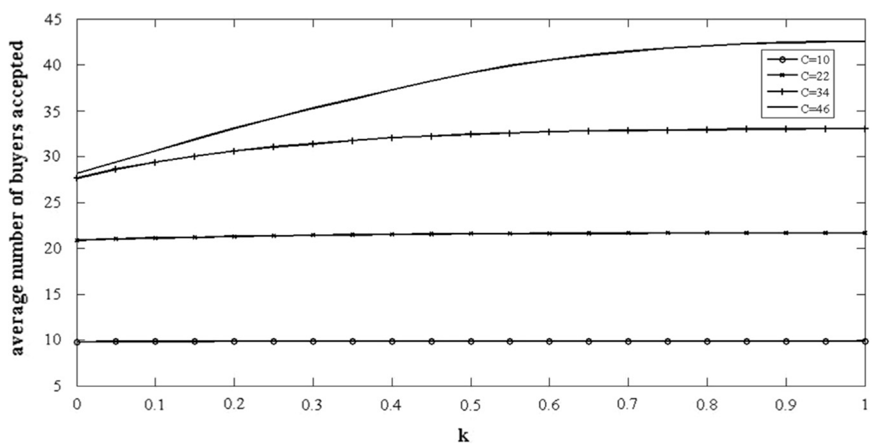

Fig. 4. Average number of successful trades across 500 instances over the k values at different C levels.

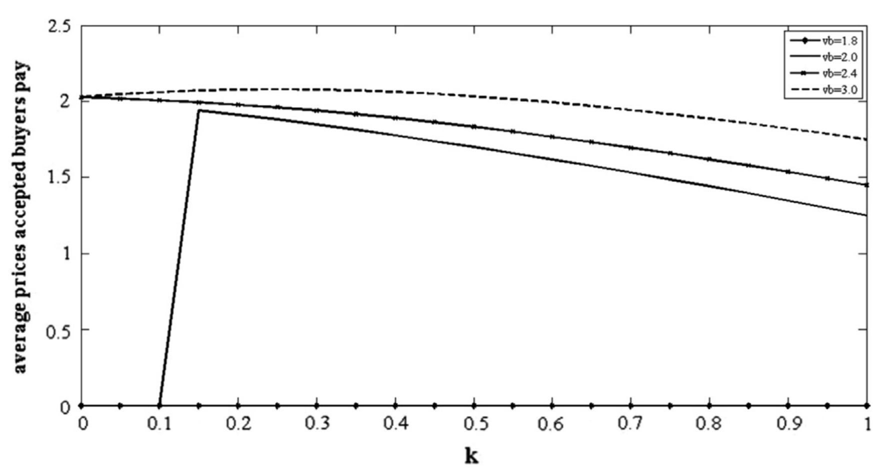

Fig. 5. : The average prices paid by buyers with different valuations across 500 instances.

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 $[0, \overline{v_b}- \underline{v_b}]$ 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 $[\underline{v_s},\overline{v_s}]=$[\$750 K, \$2000 K]; inducing $b^*(v_b)= \min\{v_b, 0.5v_b+0.5v_s\}$.

The sets of results summarized in Tables 1 and 2 illustrate the performance of the policy under uniform distribution assumption in a variety of settings as we varied the range ($[\underline{v_b},\overline{v_b}]$) the

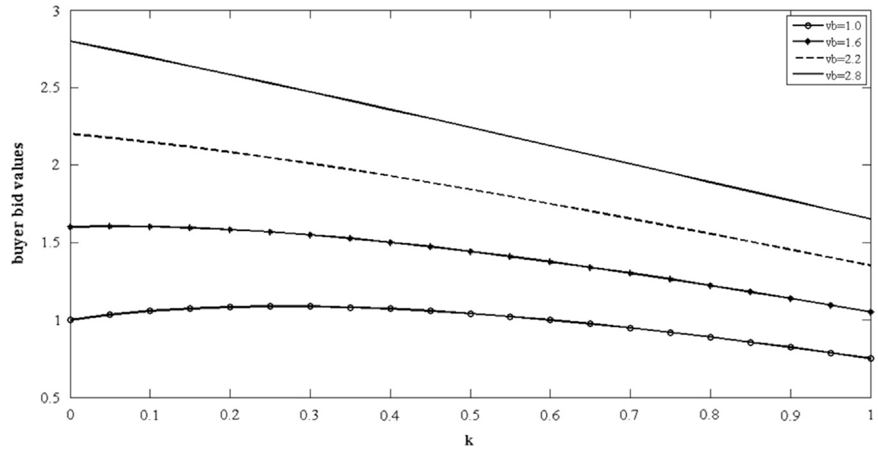

Fig. 6. : The value of buyer bids over the k values for buyers with different valuations.

Table 1
The ratio of seller's revenue under ARMC to seller's revenue under full information

	$[\underline{v_b},\overline{v_b}]$=[\$500 K, \$1500 K]C=250, C=500, C=750	$[\underline{v_b},\overline{v_b}]$=[\$1000 K, \$2000 K]C=250, C=500, C=750	$[\underline{v_b},\overline{v_b}]$=[\$1000 K, \$2500 K]C=250, C=500, C=750	$[\underline{v_b},\overline{v_b}]$=[\$1000 K, \$3000 K]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
The ratio of seller's revenue under ARMC to Seller's revenue under full information

	$[\underline{v_b},\overline{v_b}]$=[\$500 K, \$1500 K]	$[\underline{v_b},\overline{v_b}]$=[\$1000 K, \$2000 K]	$[\underline{v_b},\overline{v_b}]$=[\$1000 K, \$2500 K]	$[\underline{v_b},\overline{v_b}]$=[\$1000 K, \$3000 K]
	v_s =\$750 K, \$1000 K, \$1500 K	v_s =\$750 K, \$1000 K, \$1500 K	v_s =\$750 K, \$1000 K, \$1500 K	v_s =\$750 K, \$1000 K, \$1500 K
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%

inventory of the seller (C), and the seller valuation (v_s). In Table 1, v_s is fixed at v_s=\$1000 K, while C and $[\underline{v_b},\overline{v_b}]$ are varied to test different cases. In Table 2, C is fixed at C=500 where $[\underline{v_b},\overline{v_b}]$ 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).

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. Conclusion

In this paper, we study a monopolist seller's revenue management problem with the twist that transactions between the seller and each arriving buyer are bilaterally negotiated, a situation that has not been fully considered in the extant literature. We start with the one-to-one negotiation problems and discuss how to account for uncertainty in valuation distributions. Next, we extend our analysis to the dynamic environment: we establish the connection of the bilateral negotiation problems with the classical revenue management problems; and by studying the deterministic fluid problem, we observe the stationary nature of the optimal pricing policy. We are then able to extend the analysis to uncertain environments. Finally, two sets of numerical analyzes complement the theoretical study in other interesting perspectives, answering the questions "how the impact of parameter k in a dynamic setting might be different than in a static setting" and "how the uniform distribution assumption might affect the performance of the seller".

Although our results are limited, they offer various avenues for future research: first, several other dynamic negotiation problems may be analyzed from the perspective we presented. Of these, we believe that the games involving "strategic" buyers (or, "non-naive" buyers, as in the terminology of this paper) is of utmost interest. Another research avenue might involve characterizing the future bidding strategy of the seller with the "partial information" provided by the valuation distribution of the current buyers in a stochastic environment. Finally, the structural results regarding

Please cite this article as: Ayvaz-Cavdaroglu N, et al. Revenue management with minimax regret negotiations. Omega (2015), http://dx.doi.org/10.1016/j.omega.2015.09.009

the nature of the optimal pricing policies might be inspiring and insightful in the formulation and solution of various bilateral negotiation problems.

There are several real life settings in which a variant of the dynamic negotiation problems can be observed. These settings provide a rich scale of problems in terms of the varying negotiation power of the two parties, added uncertainty in the current and future market conditions, strategic considerations of the players, etc. We believe that there is a rich research potential in this area for interested revenue management researchers and hope that this paper serves as a starting point to lead a fruitful line of research.

Appendix A. Proof of Theorem 1

First, note that any optimal strategy should satisfy $b(v_b) \le v_b$ and $s(v_s) \ge v_s$ 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:

$$\max_{s^0}\left[(kb+(1-k)s^0-v_s).1_{\{b \ge s^0\}}-(kb+(1-k)s-v_s).1_{\{b \ge s\}}\right]$$

$$=\begin{cases} 0, & \text{if } b < v_s \\ b-v_s, & \text{if } v_s \le b \le s \\ (b-v_s)-(kb+(1-k)s-v_s). & \text{if } b > s \end{cases}$$

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:

$$\max_b \max_{s^0}\left[(kb+(1-k)s^0-v_s).1_{\{b \ge s^0\}}-(kb+(1-k)s-v_s).1_{\{b \ge s\}}\right],$$

$$=\begin{cases} 0, & \text{if } b < v_s \\ s-v_s, & \text{if } v_s \le b \le s \\ (1-k)(b-v_s), & \text{if } b > s \end{cases}$$

$$=\max\{(s-v_s),(1-k)(\overline{b}-s)\} \tag{A.1}$$

where $\overline{b}$ 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 (A.1) is increasing and the second is decreasing in s, the minimizer is attained at the intersection point, i.e:

$$s^*_{ARMC}(v_s)=\underset{s}{\operatorname{argmin}} \ \max\left\{(s-v_s),(1-k)\left(\overline{b}-s\right)\right\}$$

$$\Rightarrow s^*_{ARMC}(v_s)-v_s=(1-k)\left(\overline{b}-s^*_{ARMC}(v_s)\right)$$

$$\Rightarrow s^*_{ARMC}(v_s)=\frac{v_s}{2-k}+\frac{(1-k)}{2-k}\overline{b}$$

Via a symmetrical analysis for the buyers, we obtain $b^*_{ARMC}(v_s)=\frac{v_b}{1+k}+\frac{k}{1+k}\underline{s}$. Finally, since s^* and b^* should be best responses to each other, we find that $s^*_{ARMC}(v_s)=\frac{v_s}{2-k}+\frac{(1-k)\overline{v_b}}{2}+\frac{k(1-k)v_s}{2(2-k)}$ and $b^*_{ARMC}(v_b)=\frac{v_b}{1+k}+\frac{kv_s}{2}+\frac{k(1-k)\overline{v_b}}{2(1+k)}$. Furthermore, when the Eqs. (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 $b^*_{ARMC}(v_b)$.

Appendix B. Proof of Theorem 3

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

Since buyers are naive, their problem takes the form:

$$\underset{b}{\operatorname{argmin}}\left\{\max_s \max_{b^0}\left[\left(v_b-kb^0-(1-k)s\right).1_{\{b^0 \ge s\}}-(v_b-kb-(1-k)s).1_{\{b \ge s\}}\right]\right\} \tag{B.1}$$

$$=\max\{(v_b-b),k(b-s)\} \tag{B.2}$$

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 $s_t=s$, &for all;t, that minimizes her maximum regret for all distribution functions F_b:

$$\underset{s}{\operatorname{argmin}}\left\{\max_{F_b}\max_{s^0}\left[\int_{t=0}^{T}\Lambda_t\left[\int_{b^{-1}(s^0)}^{\overline{v_b}}(kb(v_b)+(1-k)s^0-v_s)f_b(v_b)dv_b\right]dt - \int_{t=0}^{\min\{T,T^0\}}\Lambda_t\left[\int_{b^{-1}(s)}^{\overline{v_b}}(kb(v_b)+(1-k)s-v_s)f_b(v_b)dv_b\right]dt\right]\right\}$$

where s^0 is the best bid against F_b as in (20) and T^0 is such that $\int_{t=0}^{T^0}\Lambda_t\left[\int_{b^{-1}(s)}^{\overline{v_b}}f_b(v_b)dv_b\right]dt=C$, if $s<s^0$. Regarding the inner maximization problem, we have two cases:

Case (i):$s<s^0$**:** 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. $f_b(\overline{v_b})=1$, which leads to $s^0=b(\overline{v_b})$ and the total amount of sales to be $\min\left\{C,\int_{t=0}^{T}\Lambda_t dt\right\}$. Thus:

$$\max_{F_b}\max_{s^0}\left[\int_{t=0}^{T}\Lambda_t\left[\int_{b^{-1}(s^0)}^{\overline{v_b}}(kb(v_b)+(1-k)s^0-v_s)f_b(v_b)dv_b\right]dt - \int_{t=0}^{\min\{T,T^0\}}\Lambda_t\left[\int_{b^{-1}(s)}^{\overline{v_b}}(kb(v_b)+(1-k)s-v_s)f_b(v_b)dv_b\right]dt\right]$$

$$=\min\left\{C,\int_{t=0}^{T}\Lambda_t dt\right\}\left[kb(\overline{v_b})+(1-k)b(\overline{v_b})-v_s\right] - \int_{t=0}^{\min\{T,T^0\}}\Lambda_t\left[kb(\overline{v_b})+(1-k)b(\overline{v_b})-v_s\right]dt$$

$$=\min\left\{C,\int_{t=0}^{T}\Lambda_t dt\right\}\left[(1-k)(b(\overline{v_b})-v_s)\right]$$

Case (ii):$s>s^0$**:** 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 and the seller cannot sell at all, i.e. $f_b\left(b^{-1}(s-\varepsilon)\right)=1$ for small $\varepsilon>0$. Thus, the two inner maximization problems take the form:

$$\max_{F_b}\max_{s^0}\left[\int_{t=0}^{T}\Lambda_t\left[\int_{b^{-1}(s^0)}^{\overline{v_b}}(kb(v_b)+(1-k)s^0-v_s)f_b(v_b)dv_b\right]dt - \int_{t=0}^{\min\{T,T^0\}}\Lambda_t\left[\int_{b^{-1}(s)}^{\overline{v_b}}(kb(v_b)+(1-k)s-v_s)f_b(v_b)dv_b\right]dt\right]$$

$$=\min\left\{C,\int_{t=0}^{T}\Lambda_t dt\right\}\left[k(s-\varepsilon)+(1-k)(s-\varepsilon)-v_s\right]-0$$

$$= \min \left\{C, \int_{t=0}^{T} \Lambda_t dt\right\}[s - v_s]$$

Combining the two cases, the seller should bid to minimize the two maximum regrets, i.e. $s = \text{argmin} \quad \max\left\{(s - v_s)\min\left\{C, \int_{t=0}^{T} \Lambda_t dt\right\}, (1-k)\left(\overline{b} - s\right)\min\left\{C, \int_{t=0}^{T} \Lambda_t dt\right\}\right\}$. But these two regret terms are the same as in the one-to-one game, only multiplied by a coefficient $\min\left\{C, \int_{t=0}^{T} \Lambda_t dt\right\}$. 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] Ball MO, Queyranne M. Toward robust revenue management: competitive analysis of online booking. Operations Research 2009;57(4):950–95963.

[2] Besbes O, Maglaras C. Dynamic pricing with financial milestones: feedback-form policies. Management Science 2012;58(9):1715–171731.

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

[4] Bodily S, Weatherford L. Perishable-asset revenue management: generic and multiple-price yield management with diversion. Omega – The International Journal of Management Science 1995;23(2):173–185.

[5] Campi MC, Garatti. S. The exact feasibility of randomized solutions of uncertain convex programs. SIAM Journal on Optimization 2008;19(3):1211–121230.

[6] Chatterjee K, Samuelson W. Bargaining under incomplete information. Operations Research 1983;31(5):835–83861.

[7] Chen RR, Gal-Or E, Roma P. Opaque distribution channels for competing service providers: posted price vs. name-your-own-price mechanisms. Operations Research 2014;62(4):733–73750.

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

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

[10] Feng Q, Lai G, Lu LX. Dynamic bargaining in a supply chain with asymmetric demand information. Management Science 2014;61(2):301–303.

[11] Gallego G, van Ryzin G. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. Management Science 1994;40(8):999–1020.

[12] Gallien J. Dynamic mechanism design for online commerce. Operations Research 2006;54:291–29310.

[13] Ghoniem A, Maddah B. Integrated retail decisions with multiple selling periods and customer segments: optimization and insights. Omega – The International Journal of Management Science 2015;55:38–52.

[14] Harsanyi JC. Approaches to the bargaining problem before and after the theory of games. Econometrica 1956;24(2):144–157.

[15] X. Huang, G. Soic, and G. Kersten., Selling through priceline? On the impact of name-your-own-price in competitive market. Working Paper, Marshall School of Business, University of Southern California, Los Angeles, 2014.

[16] Karmarkar US. Integrative research in marketing and operations management. Journal of Marketing Research 1996:125–133.

[17] Kim SW, Bell PC. Optimal pricing and production decisions in the presence of symmetrical and asymmetrical substitution. Omega – The International Journal of Management Science 2011;39(5):528–538.

[18] Kleindorfer PR, Wu J. Integrating long- and short-term contracting via business-to-business exchanges for capital-intensive industries. Management Science 2003;49(11):1597–1615.

[19] Kuo CW, Ahn HS, Aydin G. Dynamic pricing of limited inventories when customers negotiate. Operations research 2011;59(4):882–88897.

[20] Lan Y, Ball MO, Karaesmen I. Regret in overbooking and fare-class allocation for single leg. Manufacturing & Service Operations Management 2011;13(2):194–208.

[21] Lan Y, Ball MO, Karaesmen I, Zhang JX, Liu GX. Analysis of seat allocation and overbooking decisions with hybrid information. European Journal of Operational Research 2015;240(2):493–49504.

[22] Lan Y, Gao H, Ball MO, Karaesmen I. Revenue management with limited demand information. Management Science 2008;54(9):1594–1609.

[23] Lee TC, Hersh. M. A model for dynamic airline seat inventory control with multiple seat bookings. Transportation Science 1998;27(3):252–25265.

[24] Maglaras C, Meissner J. Dynamic pricing strategies for multi-product revenue management problems. Manufacturing & Service Operations Management 2006;8(2):136–148.

[25] Myerson RB. Two person bargaining problems with incomplete information. Econometrica 1984;2(2):461–46487.

[26] Myerson. RB. Analysis of two bargaining problems with incomplete information. Cambridge, UK: Cambridge University Press; 1985.

[27] Myerson RB, Satterthwaite MA. Efficient mechanisms for bilateral trading. Journal of Economic Theory 1983;29(2):265–281.

[28] Nash. JC. The bargaining problem. Econometrica 1950:151–162.

[29] Perakis G, Roels G. Robust controls for network revenue management. Manufacturing & Service Operations Management 2010;12(1):56–76.

[30] Phillips R, Simsek AS, van Ryzin. G. The effectiveness of field price discretion: empirical evidence from auto lending. Management Science 2015:1741–1759.

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

[32] Sen. A. A comparison of fixed and dynamic pricing policies in revenue management. Omega – The International Journal of Management Science 2012;41(3):586–597.

[33] Tasi WH, Hung SJ. Dynamic pricing and revenue management process in internet retailing under uncertainty: an integrated real options approach. Omega – The International Journal of Management Science 2009;37(2):471–47481.

[34] Terwiesch C, Savin S, Hann. IH. Online haggling at a name-your-own-price retailer: theory and application. Management Science 2005;51(3):339–351.

[35] Das Varma G, Vettas N. Optimal dynamic pricing with inventories. Economic Letters 2001;72(3):335–340.

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

[37] Zhao L, Tian P, Li X. Dynamic pricing in the presence of consumer inertia. Omega – The International Journal of Management Science 2012;40(2):137–148.

Please cite this article as: Ayvaz-Cavdaroglu N, et al. Revenue management with minimax regret negotiations. Omega (2015), http://dx.doi.org/10.1016/j.omega.2015.09.009