Pricing Optimization for Selling Reusable Resources (Extended Abstract)

Jiang Rong[1], Tao Qin[2], Bo An[3], Tie-Yan Liu[2]
[1]The Key Laboratory of Intelligent Information Processing, ICT, CAS & University of Chinese Academy of Sciences rongjiang13@mails.ucas.ac.cn
[2]Microsoft Research {taoqin, tyliu}@microsoft.com
[3]School of Computer Science and Engineering, Nanyang Technological University boan@ntu.edu.sg

ABSTRACT

The market for selling reusable products is growing rapidly. Existing works for policy optimization often ignore the dynamic property of demand and the competition among providers. This paper studies service providers' dynamic pricing in consideration of market competition and dynamics, which makes two key contributions. First, we propose a comprehensive model that takes into account the dynamic demand under market competition and formulate the optimal pricing policy as an equilibrium. Second, as it is difficult to compute the Nash equilibrium due to incomplete information and implicit revenue function, we develop an efficient algorithm to calculate an approximate equilibrium, which is more practical in the real world. The experiments show that the proposed policy outperforms existing strategies and the incentive to deviate the approximate equilibrium is small.

1. INTRODUCTION

In many real-world applications, the service providers' resources are reusable. Dynamic pricing policy plays an important role in making profits from price-sensitive users, which has shown great success in industries, e.g., the car rentals [7], hotel reservations [2, 19], network services [14], and the cloud computing [10, 20], and has attracted lots of research attention [5, 11, 15]. There are two important properties for the market: 1) users' demand is stochastic over time, which leads to dynamic inventories; and 2) providers that offer similar services need to compete against each other. However, existing works have partially neglected or treated these characteristics in an inadequate way. Against this background, this paper investigates dynamic pricing to match demand with inventory in order to maximize providers' long-term revenues in the competitive market, which gives solid theoretical and experimental analyses and makes two key contributions.

First, we propose a comprehensive model to describe the real-world applications with multiple providers and stochastic user demand, where a product can be reused, e.g., resources in a cloud platform. Existing works ignore either the competition or the dynamic feature. Demand forecast is studied in [7, 19] and the most widely-used model to describe users' dynamic demand is the Poisson process [5, 6, 14, 20]. However, those works do not consider the market competition. Xu and Hopp [22] assume that customers' arrival rates follow the geometric Brownian motion and the perfect Bayesian equilibrium is used to model providers' behaviors. Levin et al. [11] consider strategic users and propose the subgame-perfect equilibrium. However, they focus on the one-shot inventory replenishment problem with dynamic pricing, which cannot describe the market with reusable products. In this paper, we adopt the Poisson process and formulate the dynamic and competitive market as *continuous-time Markov chains* [8, 13, 16].

Since each provider aims to maximize his/her expected revenue, the optimal policy is supposed to be a *Nash Equilibrium* (NE). Our second contribution lies in that we show it is difficult to compute the NE because a provider's revenue cannot be explicitly represented as a function of his/her pricing policy and then introduce the *Approximate Equilibrium* (AE) solution concept [12, 18]. By utilizing the principles of uniformization theory [9, 17] and Bellman equation [1, 4], we propose an algorithm based on the best-response principle to efficiently compute the AE, which we demonstrate is more practical than the NE in the real market.

We conduct extensive experiments to evaluate our algorithm which shows good convergence performance. The results indicate that our pricing policy outperforms existing strategies and the incentive to deviate from the AE is tiny.

2. MODELING COMPETITIVE MARKET WITH STOCHASTIC DEMAND

We use $\mathcal{K}$ to represent the set of service providers in the market. Following the common practice in the literature [5, 14, 21, 22], we assume that users' demand for the service of provider $k \in \mathcal{K}$ is determined by two independent Poisson processes, namely the arrival process that models the coming of new demand and the departure process that corresponds to the leaving of existing requests. Specifically, we use $\lambda_k(\cdot)$ to represent the Poisson arrival rate (number of new demand instances per unit time) for provider k, which satisfies the following properties [3] $\lambda_k(p) \geq 0$, $\partial\lambda_k(p)/\partial p_k < 0$ and $\partial\lambda_k(p)/\partial p_{k'\neq k} > 0$, where $p = (p_1, p_2, \ldots, p_{|\mathcal{K}|})$. Similarly, the Poisson departure process is modeled by $\mu_k(\cdot)$, which satisfies that $\mu_k(p) \geq 0$, $\partial\mu_k(p)/\partial p_k > 0$ and $\partial\mu_k(p)/\partial p_{k'\neq k} < 0$. We use the notation $(p_k, p_{-k}) = p$. Let N_k be the maximal capacity of provider k and $[N_k]$ denote the set $\{0, 1, \ldots, N_k\}$. Since both the arrival and departure of demand are random process, the number of instances used by customers can be formulated as a continuous-time Markov process and the pricing policy of provider k is represented as $P_k = (p_{k,0}, p_{k,1}, \ldots, p_{k,N_k})$,

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

where $p_{k,n}$ is the price for the state n. Then the transition rate matrix for provider k is $Q_k(P) = (q^k_{i,j}(P))_{i,j}$, $i, j \in [N_k]$:

$$q^k_{i,j}(P) = \begin{cases} \mathbf{E}^{\pi_{-k}(P)}_{p_{-k}\in P_{-k}}\{\lambda_k(p_{k,i}, p_{-k})\}, \text{ if } j = i+1; \\ \mathbf{E}^{\pi_{-k}(P)}_{p_{-k}\in P_{-k}}\{\mu_k(p_{k,i}, p_{-k})\}, \text{ if } j = i-1; \\ -\sum_{l\neq i} q^k_{i,l}(P), \text{ if } j = i; \\ 0, \text{ otherwise}, \end{cases} \tag{1}$$

where $\pi_k(P)$ is called the *stationary* (or *steady-state*) *probability* satisfying $\sum_{n\in[N_k]} \pi_{k,n}(P) = 1$ and $\pi_k(P)\cdot Q_k(P) = 0$. When n instances are being used by customers, provider k can receive $n \cdot p_{k,n}$ revenue per unit time. Thus the average long-term expected revenue rate for provider k is

$$J_k(P_k, P_{-k}) = \sum_{n=0}^{N_k} \pi_{k,n}(P) \cdot n \cdot p_{k,n}. \tag{2}$$

3. OPTIMAL DYNAMIC PRICING

We first introduce the notation of NE.

Definition 1 (Nash equilibrium). *A Nash equilibrium is a pricing policy profile $P^* = \times_{k\in\mathcal{K}} P^*_k$, such that $\forall k \in \mathcal{K}$, $J_k(P^*_k, P^*_{-k}) \geq J_k(P_k, P^*_{-k})$ for all P_k.*

That is, no one can gain higher revenue rate by unilateral changing his/her equilibrium policy. Motivated by this observation, the NE can be computed by a *best response* procedure, which optimizes each provider j's pricing P_j while keeping others' P_{-j} fixed in each iteration, until no one wants to change his/her pricing policy. However, $J_k(P_k, P_{-k})$ is not an explicit function with respect to $P = (P_k, P_{-k})$. To address this challenge, when we optimize provider k's policy in the best-response procedure, we view the steady-state probabilities π_{-k} of others as fixed (i.e., they do not change with P). Provider k's stationary probability under this assumption, $\widehat{\pi}_k(P|\pi_{-k})$, can be calculated based on the linear equations $\sum_{n\in[N_k]} \widehat{\pi}_{k,n}(P|\pi_{-k}) = 1$ and $\widehat{\pi}_k(P|\pi_{-k}) \cdot Q_k(P|\pi_{-k}) = 0$, where $Q_k(P|\pi_{-k})$ is the same with $Q_k(P)$ except that $\pi_{-k}(P)$ in Eq.(1) is replaced with π_{-k}. The corresponding revenue rate with fixed π_{-k} is

$$\widehat{J}_k(P_k, P_{-k}|\pi_{-k}) = \sum_{n=0}^{N_k} \widehat{\pi}_{k,n}(P|\pi_{-k}) \cdot n \cdot p_{k,n}. \tag{3}$$

The optimal (best-response) policy that maximizes the above revenue rate is defined as $\widehat{B}_k(P_{-k}|\pi_{-k}) = \arg\max_{P_k\in\Delta_k} \widehat{J}_k(P_k, P_{-k}|\pi_{-k})$, which can be computed with Bellman Equation. When the random best response algorithm terminates, it follows that, $\forall k \in \mathcal{K}$, $\widehat{P}^*_k = \widehat{B}_k(\widehat{P}^*_{-k}|\pi_{-k}(\widehat{P}^*))$ and hence $\widehat{J}_k(\widehat{P}^*_k, \widehat{P}^*_{-k}|\pi_{-k}(\widehat{P}^*)) \geq \widehat{J}_k(P_k, \widehat{P}^*_{-k}|\pi_{-k}(\widehat{P}^*))$ for all $P_k \in \Delta_k$. The policy $\widehat{P}^*$ is not a NE according to Definition 1, which is an AE, as defined below.

Definition 2 (Approximate equilibrium). *An ϵ- approximate equilibrium is a pricing policy profile $\widehat{P}^* = \times_{k\in\mathcal{K}} \widehat{P}^*_k$ with a vector $\epsilon = (\epsilon_1, \epsilon_2, \ldots, \epsilon_{|\mathcal{K}|})$, such that $\forall k \in \mathcal{K}$, $J_k(\widehat{P}^*_k, \widehat{P}^*_{-k}) + \epsilon_k \geq J_k(P_k, \widehat{P}^*_{-k})$ for all P_k.*

The ϵ_k can be viewed as the additional revenue provider k can gain by unilaterally deviating from $\widehat{P}^*$, which is shown to be very small in the experiments. If $\epsilon_k = 0$ for all $k \in \mathcal{K}$, then the AE is equal to the NE. The policy $\widehat{P}^*$ is more practical than P^* in the real world with incomplete information because providers cannot calculate P^*, however, each provider k can observe others' P_{-k} and π_{-k} and then optimize his/her policy with $\widehat{B}_k(P_{-k}|\pi_{-k})$, which will make the policy to converge to the $\widehat{P}^*$ eventually.

4. EXPERIMENTAL EVALUATION

We use the following arrival and departure rate functions in our experiments: $\lambda_k(p) = l_k(1 - p_k^2)\frac{\sum_{i\neq k} p_i^2}{|\mathcal{K}|-1}$ and $\mu_k(p) = u_k p_k^2 \frac{\sum_{i\neq k}(1-p_i^2)}{|\mathcal{K}|-1}$, where l_k and u_k are parameters. To evaluate the benefits of the proposed $\widehat{P}^*$, we compare it with the existing optimal dynamic pricing [5, 14, 20], which maximizes $\sum_{k=1}^{N_k} \pi_k(P_k) n p_{k,n}$ for each provider k without consideration of others' strategy profile P_{-k}. The results are shown in Figure 1, which indicate that the noncompetitive strategy will lead to about 10% drop of revenue as compared with $\widehat{P}^*$. The evaluation for the tightness of ϵ is depicted in Table 1. We see that the benefit of deviating from $\widehat{P}^*_k$ is very limited. Thus, it is reasonable to assume providers to use $\widehat{P}^*_k$ – a more realistic equilibrium strategy that can be computed under both full and partial information assumptions.

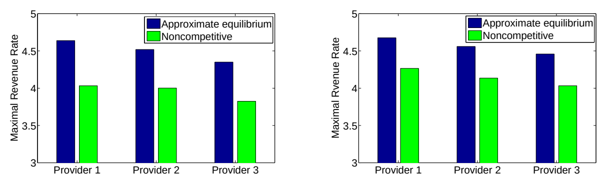

(a) S_1 (b) S_2

Figure 1: Strategy Comparison

Setting	k	l_k	u_k	N_k	$\widehat{J}^*_k(\cdot)$	ϵ_k	$\epsilon_k/\widehat{J}^*_k(\cdot)$
	1	2	1	6	4.6390	.0620	1.33%
S_1	2	1.6	1	6	4.5194	.0553	1.22%
	3	1.2	1	6	4.3506	.0463	1.06%
	1	1.6	0.8	6	4.6753	.0626	1.34%
S_2	2	1.6	1	6	4.5603	.0572	1.25%
	3	1.6	1.2	6	4.4589	.0396	0.89%

Table 1: Tightness of ϵ

5. CONCLUSION AND FUTURE WORK

We studied the dynamic pricing optimization problem for the service providers selling reusable products and made two main contributions. First, we proposed a comprehensive model that captures the dynamic and competitive features of the market. Second, we formulated providers' optimal pricing policies as an approximate equilibrium and developed an efficient algorithm to solve it. Our experimental results showed that the policy we computed outperforms existing methods in the literature. In future work, we will investigate the structural properties of the AE policy, e.g., its monotonicity with respect to the capacity utilization. Besides, we will propose more efficient algorithms to compute the best-response strategy and the AE.

Acknowledgments

This paper is supported by 2015 Microsoft Research Asia Collaborative Research Program.

REFERENCES

[1] R. E. Bellman and S. E. Dreyfus. *Applied dynamic programming*. Princeton University Press, 2015.

[2] S. Choi and A. S. Mattila. Hotel revenue management and its impact on customers' perceptions of fairness. *Journal of Revenue and Pricing Management*, 2(4):303–314, 2004.

[3] E. Dockner and S. Jørgensen. Optimal pricing strategies for new products in dynamic oligopolies. *Marketing Science*, 7(4):315–334, 1988.

[4] S. Dreyfus. Richard bellman on the birth of dynamic programming. *Operations Research*, 50(1):48–51, 2002.

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

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

[7] M. K. Geraghty and E. Johnson. Revenue management saves national car rental. *Interfaces*, 27(1):107–127, 1997.

[8] K. Gopalratnam, H. Kautz, and D. S. Weld. Extending continuous time Bayesian networks. In *Proceedings of the National Conference on Artificial Intelligence*, volume 20, page 981, 2005.

[9] D. Gross and D. R. Miller. The randomization technique as a modeling tool and solution procedure for transient Markov processes. *Operations Research*, 32(2):343–361, 1984.

[10] V. Kantere, D. Dash, G. Francois, S. Kyriakopoulou, and A. Ailamaki. Optimal service pricing for a cloud cache. *IEEE Transactions on Knowledge and Data Engineering*, 23(9):1345–1358, 2011.

[11] Y. Levin, J. McGill, and M. Nediak. Dynamic pricing in the presence of strategic consumers and oligopolistic competition. *Management Science*, 55(1):32–46, 2009.

[12] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic game theory*, volume 1. Cambridge University Press, 2007.

[13] J. R. Norris. *Markov chains*. Cambridge University Press, 1998.

[14] I. C. Paschalidis and J. N. Tsitsiklis. Congestion-dependent pricing of network services. *IEEE/ACM Transactions on Networking*, 8(2):171–184, 2000.

[15] A. Şen. A comparison of fixed and dynamic pricing policies in revenue management. *Omega*, 41(3):586–597, 2013.

[16] R. G. Simmons and H. L. Younes. Solving generalized semi-Markov decision processes using continuous phase-type distributions. In *Proceedings of the Nineteenth National Conference of Artificial Intelligence*, pages 742–747, 2004.

[17] W. J. Stewart. *Probability, Markov chains, queues, and simulation: The mathematical basis of performance modeling*. Princeton University Press, 2009.

[18] H. Tsaknakis and P. G. Spirakis. An optimization approach for approximate Nash equilibria. *Internet Mathematics*, 5(4):365–382, 2008.

[19] L. R. Weatherford and S. E. Kimes. A comparison of forecasting methods for hotel revenue management. *International Journal of Forecasting*, 19(3):401–415, 2003.

[20] H. Xu and B. Li. Maximizing revenue with dynamic cloud pricing: The infinite horizon case. In *2012 IEEE International Conference on Communications*, pages 2929–2933, 2012.

[21] H. Xu and B. Li. Dynamic cloud pricing for revenue maximization. *IEEE Transactions on Cloud Computing*, 1(2):158–171, 2013.

[22] X. Xu and W. J. Hopp. A monopolistic and oligopolistic stochastic flow revenue management model. *Operations Research*, 54(6):1098–1109, 2006.