The Hong Kong University of Science and Technology
Dept of Information Systems, Business Statistics
and Operations Management
Frontiers in Operations Management Workshop

Dynamic Pricing and Learning with Discounting

by

Dr Zhichao Feng
Assistant Professor
Department of Logistics and Maritime Studies
The Hong Kong Polytechnic University

Date	:	**2 December 2022 (Friday)**
Time	:	**4:00 – 4:45 PM**
Venue	:	**Room G012, LSK Business Building, HKUST**

Abstract:
In many practical settings, learning algorithms can take a substantial amount of time to converge, thereby raising the need to understand the role of discounting in learning. We illustrate the impact of discounting on the performance of learning algorithms by examining two classic and representative dynamic-pricing and learning problems studied in Broder and Rusmevichientong (2012) [BR] and Keskin and Zeevi (2014) [KZ]. In both settings, a seller sells a product with unlimited inventory over T periods. The seller initially does not know the parameters of the general choice model in BR (resp., the linear demand curve in KZ). Given a discount factor ρ, the retailer's objective is to determine a pricing policy to maximize the expected discounted revenue over T periods. In both settings, we establish lower bounds on the regret under any policy and show limiting bounds of $\Omega(\sqrt{1/(1-\rho)})$ and $\Omega(\sqrt{T})$ when T $\to \infty$ and $\rho \to 1$, respectively. In the model of BR with discounting, we propose an asymptotically tight learning policy and show that the regret under our policy as well that under the MLE-CYCLE policy in BR is $O(\sqrt{1/(1-\rho)})$ (resp., $O(\sqrt{T})$) when T $\to \infty$ (resp., $\rho \to 1$). In the model of KZ with discounting, we present sufficient conditions for a learning policy to guarantee asymptotic optimality, and show that the regret under any policy satisfying these conditions is $O(\log(1/(1-\rho))\sqrt{1/(1-\rho)})$ (resp., $O(\log T\sqrt{T})$) when T $\to \infty$ (resp., ρ $\to$ 1). We show that three different policies - namely, the two variants of the greedy Iterated-Least-Squares policy in KZ and a different policy that we propose - achieve this upper bound on the regret. We numerically examine the behavior of the regret under our policies as well as those in BR and KZ in the presence of discounting. We also analyze a setting in which the discount factor per period is a function of the number of decision periods in the planning horizon.

Bio:
Dr Zhichao Feng is an Assistant Professor in the Department of Logistics and Maritime Studies at The Hong Kong Polytechnic University. Before joining PolyU, he was an Assistant Professor at The University of Science and Technology of China. Zhichao obtained his Ph.D. degree from The University of Texas at Dallas and B.S. degree from Tsinghua University. His research interests focus on pricing and revenue management, queueing theory and its applications, and process analysis. His works have appeared in Management Science and Production and Operations Management.

All interested are welcome!

Enquiries: Dept of ISOM