Models of Learning in Various Games: Extended Abstract

William L. Cooper
Program in Industrial and Systems Engineering
University of Minnesota
Minneapolis, MN 55455
billcoop@me.umn.edu

Tito Homem-de-Mello
Department of Mechanical and Industrial Engineering
University of Illinois at Chicago
Chicago, IL 60607
thmello@uic.edu

Anton J. Kleywegt
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332
anton@isye.gatech.edu

December 22, 2009

Introduction Revenue management (RM) techniques typically involve the use of models by sellers to make operational pricing and availability decisions. In RM practice, settings that involve one or more sellers often have the following elements:

[A] Each seller uses a model of demand or expected revenue as a function of its own decisions (e.g., prices or booking limits).

[B] Each seller uses data comprised of its own past decisions and its own past demands (or, more accurately, sales), and occasionally of the past decisions of other sellers, to estimate the parameters of the model.

[C] With the parameter estimates in hand, each seller then treats the model and the parameter estimates as if they were correct and optimizes the objective of the model (usually expected profit or expected revenue) to make a decision.

[D] As new data are obtained, each seller updates its parameter estimates with the hope of getting better estimates and making better decisions.

We are interested in studying the following questions: How do revenue management systems behave under the interactions of elements [A]–[D] above? Do such interactions produce interesting or unexpected outcomes? Section 5.1.4.3 of Talluri and van Ryzin (2004) contains some comments on the ubiquity of monopoly models in RM practice, and their Chapter 9 discusses RM forecasting. However, there seems to be little, if any, discussion of the interaction of [A]–[D] above in the RM literature. We discuss the behavior of such dynamical systems involving one or more sellers who use models and who are attempting to use observed data to estimate the parameter values of their models.

The Duopoly Case Consider a duopoly with two sellers, called seller 1 and seller -1. Each seller i chooses a price p_i for product i, and, in response, the demand for product i is $d_i = \beta_{i,0}+\beta_{i,i}p_i+\beta_{i,-i}p_{-i}+\varepsilon_i$. The expected revenue of seller i is $g_i(p_i, p_{-i}) := p_i\left(\beta_{i,0} + \beta_{i,i}p_i + \beta_{i,-i}p_{-i}\right)$.

The Competitive Solution When the two sellers do not collude, a typical solution concept for the setting is a Nash equilibrium (NE). The unique NE prices $(p_{-1}^{\mathrm{NE}}, p_1^{\mathrm{NE}})$ are given by

$$p_i^{\mathrm{NE}} = \frac{\beta_{-i,0}\beta_{i,-i} - 2\beta_{i,0}\beta_{-i,-i}}{4\beta_{-i,-i}\beta_{i,i} - \beta_{-i,i}\beta_{i,-i}} .$$

Learning with Correct Models and Known Parameters The existence of a unique NE does not indicate whether or how such an equilibrium will arise. Suppose that sellers know the correct model and parameter values, but do not know what the others will do. A considerable variety of learning strategies has been studied for such games. (See Fudenberg and Levine 1998 for an overview.) One of the simplest such strategies is called Cournot adjustment, whereby in each period k, each seller i chooses the price that is the best response to the price chosen by seller $-i$ for the previous period $k-1$, that is,

$$p_i^k = \arg\max_{p_i} g_i(p_i, p_{-i}^{k-1}) = -\frac{\beta_{i,0} + \beta_{i,-i}p_{-i}^{k-1}}{2\beta_{i,i}} . \qquad (1)$$

It can be shown that $p_i^k \to p_i^{\mathrm{NE}}$ for $i = \pm 1$ as $k \to \infty$. It can also be shown that other types of learning rules, such as fictitious play, yield convergence to the NE as well.

Learning with Correct Models and Unknown Parameters Consider the case with two sellers in which each seller i knows the correct functional form of its demand function, but does not know the correct numerical values of the parameters. Each seller tries to learn both the numerical parameter values as well as its competitor's decision. For example, suppose that in each period $k+1$, each seller i chooses its price p_i^k using (1), but with least squares estimates in place of the unknown true parameter values. It can be shown that if the parameter estimates converge to some limits, then the prices converge to corresponding limits. However, similar to the monopoly case, it is unknown whether or not the parameter estimates do indeed converge.

Learning with an Incorrect Model Even in settings with multiple competing sellers, sellers sometimes use monopoly models, and thus these models do not explicitly account for the effects of competitors' decisions. Phillips (2005, p. 59) states "There does not appear to be a single pricing and revenue optimization system that explicitly attempts to forecast competitive response using game theory as part of its ongoing operation." It is sometimes conjectured that, although revenue management models usually do not explicitly incorporate competition, they possibly implicitly incorporate competition through parameter estimates that serve as inputs to the models. For example, see Talluri and van Ryzin (2004, p. 186) and Phillips (2005, p. 55). The question remains to what extent the effects of competition are implicitly captured in monopoly models estimated with observed data. We show that in some special settings monopoly models implicitly capture the effects of competition, but in general they do not. Another interesting question is whether sellers whose models take competition into account are better off than sellers whose models do not take competition into account. We show that it may be so, but that it is not necessarily the case. This generalizes work of Kirman (1975) and Tuinstra (2004).

References

D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press, Cambridge, MA, 1998.

A. P. Kirman. Learning by firms about demand conditions. In R. H. Day and T. Graves, editors, *Adaptive Economic Models*, pages 137–156. Academic Press, New York, 1975.

R. L. Phillips. *Pricing and Revenue Management*. Stanford University Press, Stanford, CA, 2005.

K. T. Talluri and G. J. van Ryzin. *The Theory and Practice of Revenue Management*. Kluwer Academic Publishers, Boston, 2004.

J. Tuinstra. A price adjustment process in a model of monopolistic competition. *International Game Theory Review*, 6(3):417–442, 2004.