



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

JOURNAL OF
Economic
Dynamics
& Control

Journal of Economic Dynamics & Control 27 (2003) 2207–2218

www.elsevier.com/locate/econbase

Learning competitive pricing strategies by multi-agent reinforcement learning

Erich Kutschinski^{a,*}, Thomas Uthmann^b, Daniel Polani^c

^aCentrum voor Wiskunde en Informatica, P.O. Box 94079, Amsterdam, Netherlands

^bInstitut für Informatik, Universität Mainz, Germany

^cInstitute for Neuro- and Bioinformatics, Medical University Lübeck, Germany

Abstract

In electronic marketplaces automated and dynamic pricing is becoming increasingly popular. Agents that perform this task can improve themselves by learning from past observations, possibly using reinforcement learning techniques. Co-learning of several adaptive agents against each other may lead to unforeseen results and increasingly dynamic behavior of the market. In this article we shed some light on price developments arising from a simple price adaptation strategy. Furthermore, we examine several adaptive pricing strategies and their learning behavior in a co-learning scenario with different levels of competition. Q-learning manages to learn best-reply strategies well, but is expensive to train.

© 2002 Elsevier Science B.V. All rights reserved.

JEL classification: C63

Keywords: Distributed simulation; Agent-based computational economics; Dynamic pricing; Multi-agent reinforcement learning; Q-learning

1. Introduction

For the last few years the Internet has profoundly affected the retail trade of standardized consumer goods, from books and CDs to intercontinental flights. For a single product consumers enjoy a wide choice of offers and can easily compare prices by using efficient search engines and various on-line price comparison services. But the interaction between the consumer (or the price search agent) and retailer is typically very limited—it mainly consists of obtaining price statements and eventually sending

* Corresponding author.

E-mail addresses: erich.kutschinski@cw.nl (Erich Kutschinski), uthmann@informatik.uni-mainz.de (Thomas Uthmann), polani@inb.mu-luebeck.de (Daniel Polani).

orders. For the future we envision a much more sophisticated trade on the Internet benefiting both consumers and retailers: personalized agents entering into actual negotiations would be able to act on behalf of consumers or retailers, locating specific products or variants, discussing terms of delivery or special conditions, and performing the transactions automatically, based on their owners' preferences. Market space (Eriksson et al., 1999) is one of the earliest comprehensive designs of such a system. Since then, various other commercial services allowing limited automated interaction have sprung up (for example automated bidding at ebay.com).

The currently very simple agents in these environments are in the process of being further refined to adapt their bargaining strategies to user preferences or opponent behavior. This adaptation could be effected directly through the user himself. A more efficient approach though is letting the agents learn to adapt their strategies using past experience and observations. Invariably electronic markets will become more complex and their behavior more difficult to predict, since they will consist of populations of agents adapting to each other. This scenario presents itself as a very interesting field of research.

To examine various models of such markets and their emergent behavior we developed the agent platform DMarks II at the University of Mainz (Kutschinski, 2000). It is an agent framework with decentralized control, and builds on peer-to-peer communication between its entities using the Java-RMI protocol: a design we believe is well suited to model the general structure of the large future electronic markets. The framework allows modelling the agents' strategic behavior from the bottom up, specifying precisely how an agent makes its decisions and how learning from experience is performed, and examine the resulting market quantitatively through simulations.

In this article we want to shed light on price developments in a market with just elementary adaptation rules for both buyers and sellers. Furthermore, we examine the learning capabilities of competitive seller strategies of different complexity. Different types of asynchronous multi-agent reinforcement learning (RL) will be used to determine optimal seller strategies. All experiments are set in a market scenario with an adjustable degree of competition.

The rest of the article is organized as follows: The next section provides some key facts about RL and lists related work. In Section 3 we describe the model of the market, putting more emphasis on the buyers' purchasing strategies. These are left unchanged throughout the experiments in the following sections. Section 4 discusses a market with fixed production and price developments arising from elementary seller pricing strategies. Section 5 introduces more refined pricing based on Q-learning and examines the co-learning of strategies in a market with variable supply. Finally, Section 6 summarizes our results and gives a short outlook.

2. Multi-agent RL

RL is a subfield of machine learning which addresses the task of how to learn to choose optimal actions in order to achieve a goal, given that the effect of these actions towards the goal can be perceived through sensors (Kaelbling et al., 1996).

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات