Pricing and Resource Allocation for Intelligent Trading Agents using Economic Regimes

Wolfgang Ketter⋆, John Collins, Maria Gini, Alok Gupta†, and Paul Schrater
⋆Dept. of Decision and Information Sciences, RSM Erasmus University, The Netherlands
Dept. of Computer Science and Engineering, University of Minnesota, USA
†Dept. of Information and Decision Sciences, University of Minnesota, USA
wketter@rsm.nl, {jcollins,gini,schrater}@cs.umn.edu, agupta@csom.umn.edu

Abstract

We present a semi-parametric model that describes and predicts pricing behaviors in a market environment using a Gaussian mixture model and a Markov process. We show how the model can be used to guide resource allocation and pricing decisions in an autonomous trading agent. We validate our model by presenting results obtained in the Trading Agent Competition for Supply Chain Management.

1 Introduction

We present an approach whereby an autonomous agent is able to make tactical decisions, such as product pricing, as well as strategic decisions, such as product mix and production planning, in order to maximize its expected profit in an uncertain market. The agent predicts future market conditions and adapts its decisions on procurement, production, and sales accordingly. We validate our approach on data from the Trading Agent Competition for Supply Chain Management [1](TAC SCM). TAC SCM is a market simulation in which six competing autonomous agents attempt to maximize profit by buying parts, assembling products (personal computers), and selling their products in daily auctions. The goal is to achieve the highest bank balance at the end of the game, which runs for one simulated year. An agent can produce 16 different types of computers, that are categorized into 3 different market segments. Demand in each market segment varies during the game. The paper is organized as follows. After a short overview of existing methods, we introduce our method of economic regimes for forecasting future market conditions and for modeling customer purchase behavior. In Section 4 we discuss an optimization approach to sales decisions, in particular for resource allocation and pricing. We present results of performance evaluation in Section 5. We conclude with ideas for future work.

2 Related Work

The analysis in [4] shows that the ability of decision makers to correctly identify the onset of a new regime can mean the difference between success and failure. Furthermore they found strong evidence that individuals pay inordinate attention to the signal (price in our case), and neglect diagnosticity (regime dynamics) and transition probability (Markov matrix), the aspects of the system that generates the signal. Individuals who do not pay enough attention to regime identification and prediction have the tendency to over- or underreact to market conditions. The methods which were used in the related research fail to take into account market conditions that are not directly observable. They are essentially regression models, and do not represent qualitative differences in market conditions. Our method, in contrast, is able to detect and forecast a broader range of market conditions. Regression based approaches (including non-parametric variations) assume that the functional form of the relationship between dependent and independent variables has the same structure.

An approach like ours that models variability and does not assume a functional relationship provides more flexibility and detects changes in relationship between prices and sales over time.

3 Economic Regimes and Price Prediction

Our goal is a model that can predict both future market prices and the probability that a particular bid price will result in an order. Economic theory suggests that economic environments tend to exhibit "dominant patterns" over time, such as scarcity, balanced situation, and oversupply. These patterns correspond to price ranges and trends. We call them *economic regimes*. In the following we give a brief overview of our approach. A complete description of this method is available in [2]. Our approach begins with analyzing off-line data from past sales transactions in the market. We assume that historical data are available and that they are sufficiently representative of current market conditions. Since prices are likely to have different ranges for different goods, we normalize them. We call np_g the normalized price for good g and define it as follows: $\mathrm{np_g} = price_g / nominalCost_g$ where $nominalCost_g$ is the "nominal" cost to source a unit of the good g. In TAC-SCM this is a fixed value for each good.

We estimate the price density, $p(\mathrm{np})$, by fitting a Gaussian mixture model (GMM) to historical data. We use a GMM since it is able to approximate arbitrary density functions. We use the EM-Algorithmto determine the prior probability, $P(\zeta_i)$, of the Gaussians. To obtain the price density $p(\mathrm{np})$ we sum over the products of all the individual Gaussians and their prior probability as determined by the EM algorithm. We then apply Bayes' rule to determine the posterior probability of each Gaussian given a particular normalized price observation $P(\zeta_i|\mathrm{np})$. Using Bayes' rule we transform all historical price points into posterior probability space.

We define *regimes* by clustering the price distributions across days. This is done with the k-means algorithm. The resulting clusters correspond to frequently occurring price distributions with support on a contiguous range of np. The center of each cluster is a probability vector that corresponds to regime $r = R_k$ for $k = 1, \cdots, M$, where M is the number of regimes. We compute the density of the normalized price np dependent on regime R_k from the density of the normalized price given the i-th Gaussian of the GMM and the conditional probability clustering matrix, $P(\zeta_i|R_k)$:

$$p(\mathrm{np}|R_k) = \sum_{i=1}^{N} p(\mathrm{np}|\zeta_i)\, P(\zeta_i|R_k). \tag{1}$$

The probability $P(R_k|\mathrm{np})$ can then be obtained using Bayes rule as:

$$P(R_k|\mathrm{np}) = \frac{p(\mathrm{np}|R_k)\, P(R_k)}{\sum_{k=1}^{M} p(\mathrm{np}|R_k)\, P(R_k)} \;\; \forall k = 1, \cdots, M. \tag{2}$$

where M is the number of regimes. The prior probabilities, $P(R_k)$, of the regimes are determined by a counting process over historical data. For real-time regime identification, we need to estimate the mean price of the goods sold. In TAC SCM, the agent receives every day a report that includes the minimum and maximum prices of goods sold the previous day, but not the quantities sold. The mid-range price, the price between the minimum and maximum, could be used as a coarse approximation of the mean price, but it is rather noisy. To reduce the noise we smooth yesterday's minimum and maximum prices using a Brown linear (double) exponential smoother with $\alpha = 0.5$ and compute their arithmetic mean, $\widetilde{\mathrm{np}}_{d-1}$. Since procurement, production, and pricing decisions of an agent should depend on expected future regimes, the agent needs to predict future regimes. We model the prediction of future regimes, $\vec{P}(\hat{r}_{d+h}|\widetilde{\mathrm{np}}_{d-1})$, as a Markov prediction process for strategical decision making and using exponential smoothing for tactical decision making [3].

We now describe how to predict the price trend based on regime prediction. Equation 3 describes how to compute the price density based on a predicted regime distribution. M is the number of regimes and N the number of Gaussians used in the GMM. A point on the distribution is given by

$$p(\widehat{\mathrm{np}}_{d+n}|\widetilde{\mathrm{np}}_{d-1}) \quad = \quad \sum_{j=1}^{N}\sum_{i=1}^{M} P(\zeta_j|R_i)\, P(\hat{R}_{i,d+n}|\widetilde{\mathrm{np}}_{d-1})\, p(\mathrm{np}|\zeta_j) \tag{3}$$

where $P(\hat{R}_{i,d+n}|\widetilde{\mathrm{np}}_{d-1})$ is one element of the predicted regime probability vector given either by the Markov prediction or by the exponentially smoothed predicted regimes. To obtain a predicted price distribution we sample Equation 3 for every day over the planning horizon h with values for np between 0 and 1.25[1]. After sampling the mixture distribution over the set of np values, the distribution is normalized to sum to 1. Given the estimate of the price density, the order probability function $P(order|\mathrm{np})$ can be estimated as $1 - CDF(\mathrm{np})$, where $CDF(\mathrm{np})$ is the cumulative distribution function of np.

4 Resource Allocation and Sales Pricing

Given a model for future market prices, we would like to be able to maximize the profit our agent can expect to earn over some reasonable period in the future. We do this first by allocating limited resources over time and over alternate product mix options, and second by setting prices in a way that maximizes the value of the goods we wish to sell during each sales cycle. Resource allocation decisions can be informed by experience in the past, and by observations in the present. The economic regime model encapsulates past experience and observes current market price data, giving us current and expected market prices. Other observable data:

- $\mathcal{C}$ is the set of all available component types, and $\mathcal{G}$ is the set of all goods (product types) that can be built and sold. Each good is made up of a set $\mathcal{C}_g$ of components.
- On each day d, customer demand is represented by a set $\mathcal{R}_d$ of customer RFQs received by the agent. Each RFQ $r \in \mathcal{R}_d$ specifies a product type g_r, a lead time of i_r days, a quantity q_r, and a reserve price ρ_r. Reserve price is uniformly distributed between $\rho^{\min}$ and $\rho^{\max}$. Details and semantics given in [1].
- Customer demand is projected into the future over some planning horizon h. We model customer demand with the help of a Bayesian inversion method.
- At the beginning of any given day d, the agent has an inventory of raw materials consisting of computer parts, and an inventory of finished goods consisting of different computer types.

From this data, we would like to find a way to set prices and make offers to customers that maximize the agent's overall profits. We assume that the price $price_{d,g} = f(D_{d,g}, A_{d,g})$ sustainable by the market for a given good g on a particular day d is a function of the demand $D_{d,g}$ and the quantity of goods the agent wishes to sell, the *allocation* or sales quota $A_{d,g}$ for good g on day d. The profit per unit for good g to be sold on day d at a price $price_{d,g}$ is given by $\Phi_{d,g} = discount(d)(price_{d,g} - cost(\mathcal{C}_g))$. We include the discount term as a rough approximation of inventory holding cost. It can also be used to encourage early selling, as a hedge against the uncertainty of the game. The *effective demand* function $D^{\mathit{eff}}_{d,g} = f(D_{d,g}, price_{d,g})$ for our goods will be some function of the prices $price_{d,g}$ we wish to charge. The total profit Φ over a planning horizon of h days for the set of goods $\mathcal{G}$ is then

$$\Phi = \sum_{d=0}^{h}\sum_{g\in\mathcal{G}} \Phi_{d,g} A_{d,g} \tag{4}$$

[1] 1.25 is the maximum normalized price that customers in TAC SCM are ever willing to pay.

This is what we wish to maximize, by computing values for $A_{d,g}$, subject to the following constraints:

1. We can't sell more of any product than the effective demand at the price we wish to charge.
2. For any given period of time from now until the planning horizon h, we can sell goods that we have in inventory, and goods for which we have the necessary parts in inventory.
3. The agent's factory has limited daily capacity F.

The outcome of our objective function (Eq. 4) is daily sales quotas $A_{d,g}$ for each good. The next step is to set prices so that we sell what we intend to sell, in a competitive market. In the last section, we described a model for estimating the probability of a customer placing an order as a function of price $P(order|price)$. But the quantity we sell is just the effective demand multiplied by the probability of order at the price we set. So to sell our sales quota, we need $A_{d,g} = P(order|price_{d,g})D^{eff}_{d,g}$. Once the strategic sales process has determined daily sales quotas, the next step is to set prices for our goods that will yield maximum profit. This amounts to finding, for each good, the value for $price_{d,g}$ that satisfies the relation

$$\frac{A_{d,g}}{D^{eff}_{d,g}} = P(order|price_{d,g}) \tag{5}$$

This could be solved analytically or numerically, assuming we have reasonable functions for $D^{eff}_{d,g}$ and $P(order|price_{d,g})$. In general, however, one or both of these functions are likely to be empirically-derived histograms. Under the previous assumption of most sales occurring close to a market clearing price, we can approximate $D^{eff}_{d,g}$ using $price^{est}_{d,g}$, reducing the computation to finding the value of $price_{d,g}$ that satisfies

$$\frac{A_{d,g}}{D^{eff}_{d,g}(price^{est}_{d,g})} = P(order|price_{d,g}) \tag{6}$$

The resulting customer orders provide an additional signal from the market that can be used to refine our estimate of $price^{est}_{d,g}$. If $O_{d,g}$ is the number of orders placed for good g on day d (as a result of offers made on day $d-1$), then a refined estimate of the actual market prices on that day $price^{act}_{d-1,g}$ can be found by finding an adjusted probability distribution $P^{adj}(order|price_{d-1,g})$ such that

$$\frac{O_{d,g}}{D^{eff}_{d-1,g}(price_{d-1,g})} = P^{adj}(order|price_{d-1,g}) \tag{7}$$

and computing an estimated actual price $price^{act}_{d-1,g}$ such that

$$\frac{A_{d-1,g}}{D^{eff}_{d-1,g}(price_{d-1,g})} = P^{adj}(order|price^{act}_{d-1,g}) \tag{8}$$

5 Performance evaluation

We implemented both the Markov prediction (MP) and exponential smoother (ES) prediction methods in a version of our MinneTAC TAC SCM agentand ran multiple games against a set of well-known competitors. The agents used in our experiments were obtained from the TAC SCM agent repository[2]. In addition to MinneTAC, we selected four other finalists from the 2006 competition, and an agent from the 2005 competition. For our experiments we use a controlled serverto run N_G games, each with a different pseudo-random

[2] http://www.sics.se/tac/showagents.php

sequence, with MinneTAC and the five other agents, and then run N_G games with the *same market factors* (the same set of N_G pseudo-random sequences) with a modified MinneTAC and the same set of competing agents. We use three different versions of our MinneTAC agent, each using different models for strategic decisions (price and price trend prediction) and for tactical decisions (order probability calculation). For strategic decisions we used two different price prediction methods. The first is a price-follower method (an exponential smoother predicts future prices, without using a regime model), while the second uses regimes with Markov prediction as described is Section 3 (called "Regime-M" in Table 5). For tactical decisions we used two methods to calculate the order probability. The first is a simple linear interpolation between the smoothed minimum and maximum prices, the second uses the regime model and makes predictions using the exponential smoother (called "Regime-E" in Table 1.

Experiment **Strategic:** **Tactical:**	1 Follower Linear	2 Regime-M Linear	3 Regime-M Regime-E
Agent	Mean Profit/Std. Dev. (in $M)		
TacTex	8.75/5.68	8.87/5.60	9.21/5.39
DeepMaize	8.84/4.63	8.71/4.85	8.32/4.18
PhantAgent	8.05/5.42	7.99/5.38	8.17/5.44
Maxon	4.24/4.52	3.77/4.29	4.02/4.18
MinneTAC	1.35/3.70	1.81/4.02	2.12/3.76
Rational	0.74/4.91	0.67/4.69	1.31/4.53

Table 1: Experimental results with repeated market conditions and three variations of MinneTAC for order probability, price and price trend predictions. Mean profit and standard deviation results are based on 23 games. Regime-M uses the regime model with Markov prediction process, and Regime-E uses the regime model with exponential smoother lookup process.

6 Conclusions and Future Work

We have presented an approach for strategic and tactical decision making in a competitive sales environment, based on predicting market prices and price trends, optimizing product mix and resource allocation, and estimating the probability of receiving an order for a given offer price. Currently we are working on weighted regime ensemble prediction, which means that the different predictions models get rewarded depending on their performance. We intend to apply our method in other domains such as Amazon.com, eBay.com, and in financial applications like stock tracking and forecasting.

References

[1] John Collins, Raghu Arunachalam, Norman Sadeh, Joakim Ericsson, Niclas Finne, and Sverker Janson. The Supply Chain Management Game for the 2006 Trading Agent Competition. Technical Report CMU-ISRI-05-132, Carnegie Mellon University, Pittsburgh, PA, November 2005.

[2] Wolfgang Ketter. *Identification and Prediction of Economic Regimes to Guide Decision Making in Multi-Agent Marketplaces*. PhD thesis, University of Minnesota, Twin-Cities, USA, January 2007.

[3] Wolfgang Ketter, John Collins, Maria Gini, Alok Gupta, and Paul Schrater. Detecting and Forecasting Economic Regimes in Automated Exchanges. *Decision Support Systems*, page Forthcoming, 2007.

[4] Cade Massey and George Wu. Detecting regime shifts: The causes of under- and overestimation. *Management Science*, 51(6):932–947, 2005.