Strategic Sales Management Guided By Economic Regimes

Wolfgang Ketter, John Collins, Maria Gini*, Alok Gupta†, and Paul Schrater

Department of Computer Science and Engineering
†Department of Information and Decision Sciences
University of Minnesota, Minneapolis, MN 55455, USA
{ketter,jcollins,gini,schrater}@cs.umn.edu, agupta@csom.umn.edu

Abstract. We present methods to characterize market conditions from historical data, and we describe how this knowledge can be used to make strategic and tactical sales decisions. The methods are based on learning dominant market conditions, such as over-supply or scarcity, from historical data using computational methods to represent the price density function. We show how to use this knowledge, together with real-time observable information, to identify the current dominant market condition and to forecast market changes over a planning horizon. We validate our methods by presenting experimental results in a case study, the Trading Agent Competition for Supply Chain Management.

1 Introduction

Business organizations have an increasing need for software that can assist decision makers by gathering and analyzing information and making recommendations and business decisions. Advanced decision support systems and autonomous software agents promise to address this need by acting rationally on behalf of humans in numerous application domains. Examples include procurement [1, 2], scheduling and resource management [3, 4], and personal information management [5, 6].

The approach we present in this paper is based on using machine learning algorithms to support rational decision making by an autonomous agent that operates in a supply chain environment. In our method we characterize market conditions by distinguishable statistical patterns, that we call *regimes*. We show how such patterns can be learned from historical data and identified from observable data. We outline how to identify regimes and forecast regime transitions. This prediction, in turn, can be used to allocate resources to current and future sales in a way that maximizes resource value. While this type of prediction about the economic environment is commonly used at the macro economic level [7], such predictions are rarely done for a micro economic environment.

We use a competitive trading environment, the Trading Agent Competition for Supply Chain Management [8] (TAC SCM) as a testbed for our methodology.

* Partial funding is acknowledged from NSF under grant IIS-0414466.

TAC SCM is a market simulation in which six autonomous agents buy parts, assemble personal computers, and sell them to consumers in daily auctions.

A TAC SCM agent must base its decision on how to allocate its resources and how to set prices using limited information about the state of the market and the strategies of the other agents. Agents must simultaneously compete in two separate but interrelated markets: the market from which the agents must buy their supplies and the market to which the agents must sell their finished products. Agents have a large number of decisions to make in a limited time, so computational efficiency of the decision-making process is essential. Therefore, TAC SCM is an appropriate and comprehensive testbed for our methodology.

Although in this paper we present regime identification, prediction, and resource allocation in the context of TAC SCM, our method is applicable to other domains, both for autonomous decision making by an agent and as a decision support technology. Examples of domains to which our proposed approach can be applied include agents for automated trading in financial markets, such as the Penn-Lehman Automated Trading Project [9], auction-based contracting environments, such as MAGNET [10], and other auctions, such as auctions for IBM PCs [11] or PDA's on eBay [12].

After a review of relevant literature, we describe in a general way the information needed to make strategic and tactical sales decisions. We follow with a discussion of the concept of "economic regimes" and their representation using learned probability density functions. We then describe how this method is used in an automated trading agent. For reader's convenience, we present a summary of our notation in the Appendix.

2 Related Literature

Massey and Wu [13] show in their analysis 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 probabilities) 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.

In [14] the authors empirically analyze the degree to which used products cannibalize new product sales for books on Amazon.com. In their study they show that product prices go through different regimes over time. Marketing research methods have been developed to understand the conditions for growth in performance and the role that marketing actions can play to improve sales. For instance, in [15], an analysis is presented on how in mature economic markets strategic windows of change alternate with long periods of stability.

Much work has focused on models for rational decision-making in autonomous agents. Ng and Russel [16] show that an agent's decisions can be viewed as a set of linear constraints on the space of possible utility (reward) functions. However, the simple reward structure they used in their experiments will not scale to what is needed to predict prices in more complex situations such as TAC SCM.

Sales strategies used in previous TAC SCM competitions have attempted to model the probability of receiving an order for a given offer price, either by estimating the probability by linear interpolation from the minimum and maximum daily prices [17], or by estimating the relationship between offer price and order probability with a linear cumulative density function (CDF) [18], or by using a reverse CDF and factors such as quantity and due date [19].

All these methods 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.

An analysis [20] of the TAC SCM 2004 competition shows that supply and demand (expressed as regimes in our method) are key factors in determining market prices, and that agents which were able to detect and exploit these conditions had an advantage.

3 Sales Strategy

We assume that a rational agent will attempt to maximize its profit. In an environment where production capacity and component supplies are limited, sales must be focused on getting the highest possible prices for the resources the agent has available over some planning horizon. Suboptimal sales decisions may result in a situation where inventory is sold out during a period when prices are low, leaving the agent with nothing available to sell later when prices are higher.

Typically, an agent makes sales decisions in two steps. The first step is a strategic decision, where resources are allocated over a planning horizon in a way that maximizes expected profit. The second step is a tactical decision, which determines the prices that are expected to sell the quantities determined by the strategic decision, given the current customer demand and pricing model. In this paper we focus primarily on the strategic decisions.

Figure 1 gives a high-level architecture of the components involved in making sales decisions. As illustrated, some data are expected to be external to the agent and observable in the market. This includes things such as daily price reports, data on customer demand, and data on sales made recently. Some information, such as inventory status, resource constraints, and cost basis, is expected to be internal to the agent and not directly visible to any other agent. Some information is expected to be derived by the agent using its own prediction methods.

The two major components that are used for decision making are the “Economic Regime Model” and the “Allocation” modules. The “Economic Regime

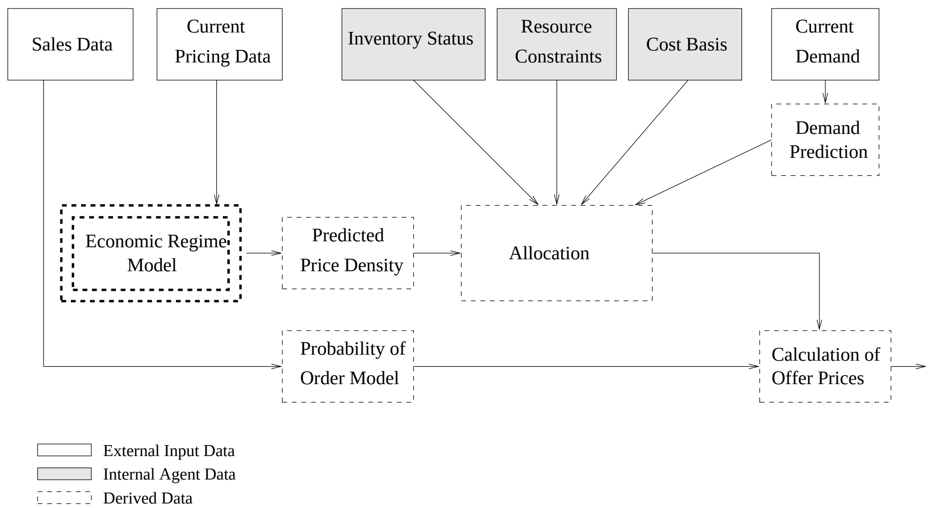

Fig. 1. Pricing Chain. Economic regime modeling (double borders) is the focus of this paper. It is a tool for tactical and strategic decision making. The calculation of the offer prices by the Allocation module is a tactical decision.

Model" uses current and past sales data to predict regimes and price density of the goods to be sold over some time horizon. The "Allocation" module sets the quantity (sales quota) the agent wishes to sell through some time horizon. The decision is made taking into account resource constraints, current inventory status, cost basis, and expected demand. This can be done, for example, by solving a linear program that maximizes the total profit over the selected horizon and over the set of finished goods that the agent can potentially sell. If the agent can predict the approximate market price for each product to be sold in the planning horizon, then it can compute the predicted profit per unit of good.

4 Economic Regimes

Market conditions change over time, and this should affect the strategy used by an agent in procurement, production planning, and product pricing. Economic theory suggests that economic environments exhibit 3 dominant market patterns: scarcity, balanced, and over-supply. We define a scarcity condition if there is more customer demand than product supply in the market, a balanced condition if demand is approximately equal to supply, and an over-supply condition if there is less customer demand than product supply in the market. When there is scarcity, prices are higher, so the agent should price more aggressively. In balanced situations, prices are lower and have more spread, so the agent has a range of options for maximizing expected profit. In over-supply situations prices are lower. The agent should primarily control costs, and therefore either do pricing based on costs, or wait for better market conditions.

We believe that even though the market is constantly changing, there are some underlying dominant patterns which characterize the aforementioned market conditions. We define a specific mode a market can be in as a *regime*. A way of solving the decision problem an agent is faced with is to characterize those regimes and to apply specific decision making methods to each regime. This requires an agent to have methods for figuring out what is the current regime and for predicting which future regimes will be in its planning horizon.

4.1 Analysis of historical data to characterize market regimes

The first phase in our approach is to identify and characterize market regimes by analyzing data from past sales. The assumption we make is that enough historical data are available for the analysis and that historical data are sufficiently representative of possible market conditions. Information observable in real-time in the market is then used to identify the current regime and to forecast regime transitions.

Since product prices are likely to have different ranges for different products, we normalize them. We call np the normalized price and define it as follows:

$$\begin{aligned}\mathrm{np} &= \frac{ProductPrice}{NominalProductCost} \\ &= \frac{ProductPrice}{AssemblyCost + \sum_{j=1}^{j=numParts} NominalPartCost_j}\end{aligned}$$

where $NominalPartCost_j$ is the nominal cost of the j-th part, $numParts$ is the number of parts needed to make the product, and $AssemblyCost$ is the cost of manufacturing the product. An advantage of using normalized prices is that we can easily compare price patterns across different products.

Historical data are used to estimate the price density, $p(\mathrm{np})$, and to characterize regimes. We start by estimating the price density function by fitting a Gaussian mixture model (GMM) to historical normalized price, np, data. We use a GMM since it is able to approximate arbitrary density functions. Another advantage is that the GMM is a semi-parametric approach which allows for fast computing and uses less memory than other approaches. We use the Expectation-Maximization (EM) Algorithm [21] to determine the Gaussian components of the GMM, $N[\mu_i, \sigma_i](\mathrm{np})$, and their prior probabilities, $P(c_i)$. The density of the normalized price can be written as:

$$p(\mathrm{np}) = \sum_{i=1}^{N} p(\mathrm{np}|c_i)\, P(c_i)$$

where $p(\mathrm{np}|c_i)$ is the i-th Gaussian from the GMM, i.e.,

$$p(\mathrm{np}|c_i) = p(\mathrm{np}|\mu_i, \sigma_i) = \frac{1}{\sigma_i\sqrt{2\pi}} e^{\left[\frac{-(\mathrm{np}-\mu_i)^2}{2\times\sigma_i^2}\right]}$$

where μ_i is the mean and σ_i is the standard deviation of the i-th Gaussian from the GMM. An example of a GMM is shown in Figure 2. While the choice of N, the number of Gaussians', in a GMM is arbitrary, the choice should reflect a balance in derived accuracy and computational overhead. We chose $N = 10$, because we found experimentally that this provides good quality. We tried other values, ranging from 5 to 15, with similar results.

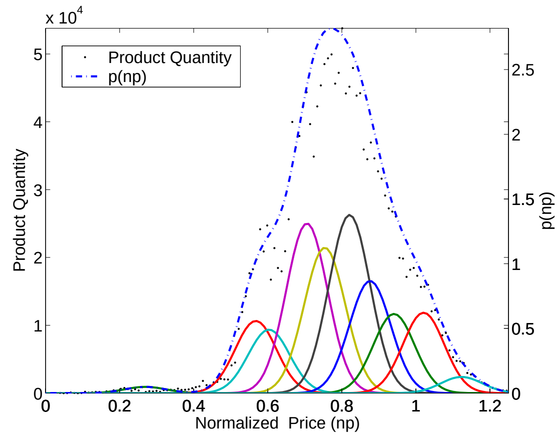

Fig. 2. The price density density function, $p(np)$, (right y-axis) estimated by the Gaussian mixture model fits well the historical normalized price data (left y-axis represents product quantity) for a sample market. Data are from 18 games from semi-finals and finals of TAC SCM 2005.

Using Bayes' rule we determine the posterior probability of the i-th Gaussian of the GMM given a normalized price, np, as follows:

$$P(c_i|\mathrm{np}) = \frac{p(\mathrm{np}|c_i)\, P(c_i)}{\sum_{i=1}^{N} p(\mathrm{np}|c_i)\, P(c_i)} \qquad \forall i = 1, \cdots, N$$

We then define the posterior probabilities of all Gaussians' given a normalized price, np, as the following N-dimensional vector:

$$\boldsymbol{\eta}(\mathrm{np}) = [P(c_1|\mathrm{np}), P(c_2|\mathrm{np}), \ldots, P(c_N|\mathrm{np})].$$

For each normalized price np_j we compute the vector of the posterior normalized price probabilities, $\boldsymbol{\eta}(\mathrm{np}_j)$, which is $\boldsymbol{\eta}$ evaluated at each observed normalized price np_j.

We cluster these collections of vectors using the k-means algorithm. Each cluster corresponds to a $r$egime. 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. Collecting these vectors into a matrix yields the conditional probability

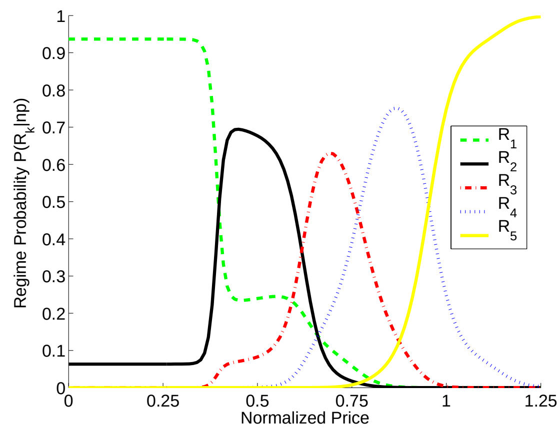

Fig. 3. An example of learned regime probabilities , $P(R_k|\mathrm{np})$, over normalized price np, for a sample market in TAC SCM after training.

matrix $\mathbf{P}(c|r)$. The matrix has N rows, one for each component of the GMM, and M columns, one for each regime.

In Figure 3 we distinguish five regimes, which we can call extreme over-supply (R_1), over-supply (R_2), balanced (R_3), scarcity (R_4), and extreme scarcity (R_5). Regimes R_1 and R_2 represent a situation where there is a glut in the market, i.e. an over-supply situation, which depresses prices. Regime R_3 represents a balanced market situation, where most of the demand is satisfied. In regime R_3 the agent has a range of options of price vs sales volume. Regimes R_4 and R_5 represent a situation where there is scarcity of products in the market, which increases prices. In this case the agent should price as close as possible to the estimated maximum price a customer is willing to pay.

For the TAC SCM domain, the number of regimes was selected a priori, after examining the data and looking at economic analyses of market situations. In our experiments we found out that the number of regimes chosen does not significantly affect the results regarding price trend predictions. The computation of the GMM and k-means clustering were tried with different initial conditions, but consistently converged to the same results.

We marginalize the product of the density of the normalized price, np, given the i-th Gaussian of the GMM, $p(\mathrm{np}|c_i)$, and the conditional probability clustering matrix, $P(c_i|R_k)$, over all Gaussians c_i. We obtain the density of the normalized price np dependent on the regime R_k:

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

The probability of regime R_k dependent on the normalized price np can be computed using Bayes rule as:

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

where M is the number of regimes. The prior probabilities, $P(R_k)$, of the different regimes are determined by a counting process over past data. Figure 3 depicts the regime probabilities for a sample market in TAC SCM. Each regime is clearly dominant over a range of normalized prices.

The intuition behind regimes is that prices communicate information about future expectations of the market. However, absolute prices do not mean much because the same price point can be achieved in a static mode (i.e., when prices don't change), when prices are increasing, or when prices are decreasing. In the construction of a regime the variation in prices (the nature, variance, and the neighborhood) are considered thereby providing a better assessment of market conditions.

The last step is the computation of a Markov transition matrix to be used by the agent for regime prediction. We model regime prediction as a Markov process. We construct a Markov transition matrix, $\mathbf{T}_{\mathrm{predict}}(r_{t+1}|r_t)$ by a counting process over past data. This matrix represents the posterior probability of transitioning at time step $t+1$ to regime r_{t+1} given the current regime r_t at time step t.

4.2 Identification of current regime

Previous sales data are used to learn the characterization of different market regimes. In real-time an agent can then use this regime information to identify the dominant regime. This can be done by calculating (or estimating) the normalized prices for the current time step, t.

Since complete current price information might not be available, we indicate the estimated normalized price at time t by $\overline{\mathrm{np}}_t$. Depending on the application domain, the price estimate can be accurate, or can be an approximation.

The agent selects the regime which has the highest probability, i.e.

$$\underset{1 \le k \le M}{\operatorname{argmax}} \boldsymbol{P}(R_k|\overline{\mathrm{np}}_t).$$

4.3 Regime prediction

The prediction of regime probabilities is based on two distinct operations:

1. a correction (recursive Bayesian update) of the posterior probabilities for the regimes based on the history of measurements of the estimated median normalized prices obtained since the time of the last regime change until the previous time step.
2. a prediction of regime posterior probabilities for the current time step. The prediction of the posterior distribution of regimes n time steps into the future is done recursively.

The agent can use the regime identification and the forecast of regime transitions to adapt its procurement, production, and pricing strategies accordingly.

5 A Case Study: TAC SCM

The Trading Agent Competition for Supply Chain Management [8] (TAC SCM) is a market simulation in which six autonomous agents compete to maximize profits in a computer-assembly scenario. The simulation takes place over 220 virtual days, each lasting fifteen seconds of real time. Agents earn money by selling computers they assemble out of parts they purchase from suppliers. Each agent has a bank account with an initial balance of zero. The agent with the highest bank balance at the end of the game wins.

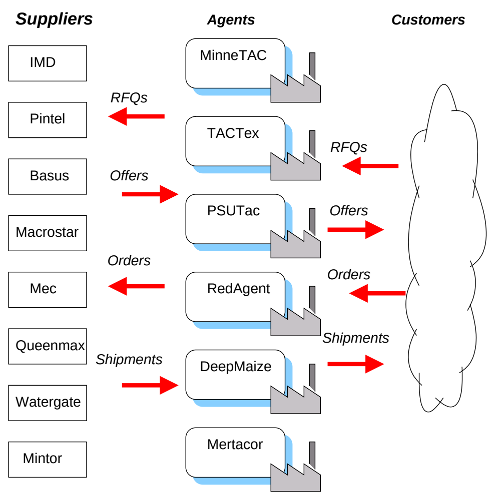

Fig. 4. Schematic overview of a typical TAC SCM game scenario.

To obtain parts, an agent must send a *request for quotes* (RFQ) to an appropriate *supplier*. Each RFQ specifies a component type, a quantity, and a due date. The next day, the agent will receive a response to each request. Suppliers respond by evaluating each RFQ to determine how many components they can deliver on the requested due date, considering the outstanding orders they have committed to and at what price. If the supplier can produce the desired quantity on time, it responds with an offer that contains the price of the supplies. If not, the supplier responds with two offers: (1) an earliest complete offer with a revised due date and a price. This revised due date is the first day in which the supplier believes it will be able to provide the entire quantity requested; and (2) a partial offer with a revised quantity and a price with the requested due date. The agent can accept either of these alternative offers, or reject both. Suppliers may deliver late, due to uncertainty in their production capacities. Suppliers discount part prices according to the ratio of supply to demand.

Every day each agent receives a set of RFQs from potential *customers*. Each customer RFQ specifies the type of computers requested, along with quantity, due date, reserve price, and penalty for late delivery. Each agent may choose to bid on some or all of the day's RFQs. Customers accept the lowest bid that is at or below the reserve price, and notify the winning agent. The agent must ship customer orders on time, or pay the penalty for each day an order is late. If a product is not shipped within five days of the due date the order is canceled, the agent receives no payment, and no further penalties accrue.

An agent can produce 16 different types of computers, that are categorized into three different market segments (low, medium, and high). Demand in each market segment varies randomly during the game. Other variables, such as storage costs and interest rates also vary between games.

The other agents playing in the same game affect significantly the market, since they all compete for the same parts and customers. This complicates the operational and strategic decisions an agent has to make every day during the game, which include how many parts to buy, when to get the parts delivered, how to schedule its factory production, what types of computers to build, when to sell them, and at what price.

5.1 Experimental setup

For our experiments, we used data from a set of 24 games (18 for training[1] and 6 for testing[2]) played during the semi-finals and finals of TAC SCM 2005. The mix of players changed from game to game, the total number of players was 12 in the semi-finals and 6 in the finals.

Since supply and demand in TAC SCM change in each of the market segments (low, medium, and high) independently of the other segments, our method is applied to each individual market segment.

Each type of computer has a nominal cost, which is the sum of the nominal cost of each of the parts needed to build it. In TAC SCM the cost of the facility is sunk, and there is no per-unit assembly cost. We normalize the prices across the different computer types in each market segment.

5.2 Online identification of current regime

Every day the agent receives a report which includes the minimum and maximum prices of the computers sold the day before, but not the quantities sold. We use the mid-range price, $\overline{\text{np}}$, the price between the minimum and maximum, to

[1] 3694@tac3, 3700@tac3, 4229@tac4, 4234@tac4, 7815@tac5, 7821@tac5, 5638@tac6, 5639@tac6, 3719@tac3, 3720@tac3, 3721@tac3, 3722@tac3, 3723@tac3, 4255@tac4,4256@tac4, 4257@tac4, 4258@tac4, 4259@tac4 – To obtain the complete path name append .sics.se to each game number.

[2] 3717@tac3, 3718@tac3, 3724@tac3, 4253@tac4, 4254@tac4, 4260@tac4

approximate the mean price. This approximation sometimes is quite poor. An example which shows how the mid-range value differs from the mean value is in Figure 5. The mean value is computed after the game, when the entire game data are available. We observe that the mid-range price is different from the mean price. In this example, around day 40 and occasionally beyond day 120, we observe a high spike in the maximum price. This was caused by an opportunistic agent who discovered a small amount of unsatisfied demand, but most of that day's orders were sold at a much lower price. To get a better approximation, we apply a double exponential smoother on the mid-range price. Figure 5 shows that the smoothed mid-range price, $\tilde{\mathrm{np}}$, is closer to the mean price.

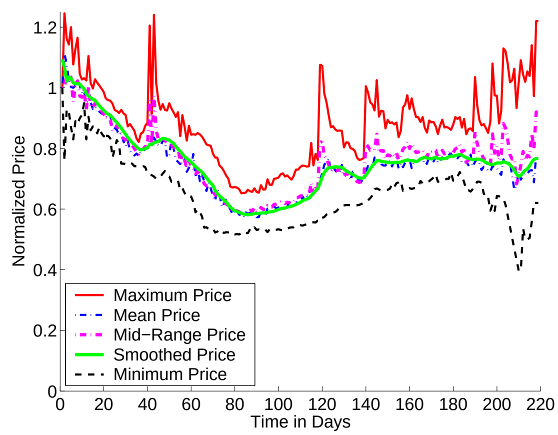

Fig. 5. Minimum, maximum, mean, mid-range, and smoothed mid-range daily normalized prices of computers sold. All prices except the mean are computed from the data reported during the game every day. The mean price is computed after the game using the game data, which include complete information on all the transactions. Data are for the high market segment in 4257@tac4, one of the final games.

During the game, the agent estimates on day t the current regime by calculating the smoothed mid-range normalized price $\tilde{\mathrm{np}}_{t-1}$ for the previous day (recall that the agent every day receives the minimum and maximum prices for the previous day) and by selecting the regime which has the highest probability, i.e. $\mathrm{argmax}_{1\leq k\leq M}\, \boldsymbol{P}(R_k|\tilde{\mathrm{np}}_{t-1})$.

Figure 6 shows the relative probabilities of each regime over the course of a game. The graph shows that different regimes are dominant at different points in the game, and that there are brief intervals during which two regimes are almost equally likely. An agent can use this information to decide which strategy, or mixture of strategies, to follow. When making decisions the agent should consider not only the current regime (tactical decision), but also upcoming regime changes (strategic decision). Therefore an agent needs to predict future regimes.

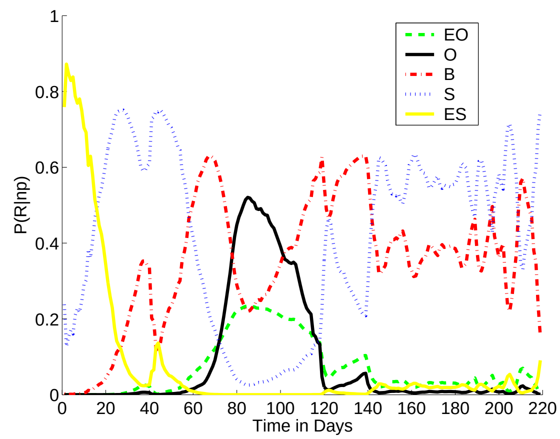

Fig. 6. Regime probabilities over time computed online every day for the high market segment. Data are from game 4257@tac4.

5.3 Regime prediction

Since in the context of TAC SCM each time step corresponds to a single day we train the Markov matrix, $\mathbf{T}_{\text{predict}}(r_{t+1}|r_t)$, on regime changes from day t to day $t+1$ by a counting process over past games. The correction and prediction steps are applied to the current regime probabilities as outlined in Section 4.3. Examples of regime predictions for a sample game for the high market segment are shown in Figure 7 and Figure 8.

The figures show the real regimes measured after the game from the game data and the predictions made by our method during the game. As it can be seen in the figures, the match between predictions and real data is very good.

Figure 7 shows a predicted change from extreme-scarcity to scarcity. In this case the agent should try to sell more aggressively the current day, since prices will be decreasing in the next days. On the other hand we see in Figure 8 a change from an extreme-oversupply situation to an oversupply situation. This means that the agent should sell less today and build up more inventory for future days when prices will be higher.

We have reported in [22] and [23] a correlation analysis of the market parameters to regimes and more details on regime identification and prediction.

5.4 Tactical Sales Decision

The regime model we have described can help the agent to make strategic decisions that can drive procurement, production, and sales behaviors. We now briefly outline how our agent makes tactical sales decisions.

Given the daily sales quotas, price trend predictions, and the current demand for each product, the final tactical decision is to set the offer price for each product to a value that is expected to sell the desired quota. We assume that

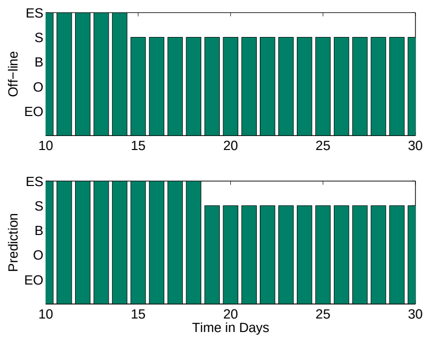

Fig. 7. Regime predictions for game 4257@tac4 starting on day 10 for 20 days into the future for the high market segment. Data are shown as computed after the game using the complete set of data, and as predicted by our method during the game.

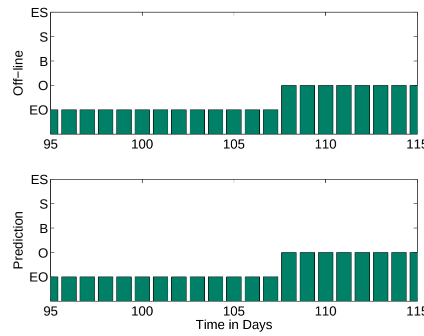

Fig. 8. Regime predictions for game 4257@tac4 starting on day 95 for 20 days into the future for the high market segment.

the computed price is offered for the entire demand. This is done with the help of a pricing model that approximates the probability that customers will accept an offer at a given price. The agent uses a simple linear approximation, given the expected median price (derived from $\widetilde{np}$ and from a feedback loop that monitors the agent's sales performance) and slope of the acceptance probability function. Offer prices are slightly randomized, and actual sales performance is used to update the model.

6 Conclusions and Future Work

We have presented an approach for identifying and predicting market conditions in markets for durable goods. We have demonstrated the effectiveness of our approach using games played in the semi-finals and finals from TAC SCM 2005. An advantage of the proposed method is that it works in any market for durable goods, since the computational process is completely data driven and that no classification of the market structure (monopoly vs competitive, etc) is needed.

Our approach recognizes that different market situations have qualitative differences that can be used to guide the strategic and tactical behavior of an agent. Unlike regression-based methods that try to predict prices directly from demand and other observable factors, our approach recognizes that prices are also influenced by non-observable factors, such as the inventory positions of the other agents.

Our method promises to enable an agent to anticipate and prepare for regime changes, for example by building up inventory in anticipation of better prices in the future or by selling in anticipation of an upcoming oversupply situation. In future work we plan to apply our method to data from Amazon.com or eBay.com.

7 Appendix: Summary of Notation

Symbol	Definition
np	Normalized price
$\overline{\mathrm{np}}$	Mid-range normalized price
$\tilde{\mathrm{np}}$	Smoothed mid-range normalized price
$p(\mathrm{np})$	Density of the normalized price
GMM	Gaussian Mixture Model
N	Number of Gaussians of the GMM
$p(\mathrm{np}\|c_i)$	Density of the normalized price, np, given i-th Gaussian of the GMM
μ_i	Mean of i-th Gaussian of the GMM
σ_i	Standard deviation of i-th Gaussian of the GMM
$P(c_i)$	Prior probability of i-th Gaussian of the GMM
$P(c_i\|\mathrm{np})$	Posterior probability of the i-th Gaussian of the GMM given a normalized price np
$\boldsymbol{\eta}(\mathrm{np})$	N-dimensional vector with posterior probabilities, $P(c_i\|\mathrm{np})$, of the GMM
M	Number of regimes
R_k	k-th Regime, $k = 1, \cdots, M$
$\mathbf{P}(c\|r)$	Conditional probability matrix (N rows and M columns) resulting from k-means clustering
$p(\mathrm{np}\|R_k)$	Density of the normalized price np dependent on the regime R_k
$P(R_k\|\mathrm{np})$	Probability of regime R_k dependent on the normalized price np
t	Current time
t_0	Time of last regime change
$\mathbf{T}_{\mathrm{predict}}(r_{t+1}\|r_t)$	Markov transition matrix

References

1. Sandholm, T.: Expressive commerce and its application to sourcing. In: Proc. of the Twenty-First National Conference on Artificial Intelligence, Boston, MA, AAAI (2006) 1736–1743
2. CombineNet: Sourcing solutions. http://www.combinenet.com/sourcing_solutions/ (2006)
3. I2: Next-generation planning. http://i2.com/solution_library/ng_planning.cfm (2006)
4. Collins, J., Bilot, C., Gini, M., Mobasher, B.: Decision processes in agent-based automated contracting. IEEE Internet Computing **5**(2) (2001) 61–72
5. Berry, P., Conley, K., Gervasio, M., Peintner, B., Uribe, T., Yorke-Smith, N.: Deploying a personalized time management agent. In: Proc. of the Fifth Int'l Conf. on Autonomous Agents and Multi-Agent Systems, Hakodate, Japan (2006)
6. Mark, B., Perrault, R.C.: Calo: Cognitive assistant that learns and organizes. http://www.ai.sri.com/project/CALO (2006)

7. Osborn, D.R., Sensier, M.: The prediction of business cycle phases: financial variables and international linkages. National Institute Econ. Rev. **182**(1) (2002) 96–105
8. Collins, J., Arunachalam, R., Sadeh, N., Ericsson, J., Finne, N., Janson, S.: The Supply Chain Management Game for the 2005 Trading Agent Competition. Technical Report CMU-ISRI-04-139, Carnegie Mellon University, Pittsburgh, PA (2004)
9. Kearns, M., Ortiz, L.: The Penn-Lehman Automated Trading Project. IEEE Intelligent Systems (2003) 22–31
10. Collins, J., Ketter, W., Gini, M.: A multi-agent negotiation testbed for contracting tasks with temporal and precedence constraints. Int'l Journal of Electronic Commerce **7**(1) (2002) 35–57
11. Lawrence, R.: A machine learning approach to optimal bid pricing. In: 8th INFORMS Computing Society Conf. on Optimization and Computation in the Network Era, Arizona (2003)
12. Ghani, R.: Price prediction and insurance for online auctions. In: Int'l Conf. on Knowledge Discovery in Data Mining, Chicago, Illinois (2005) 411–418
13. Massey, C., Wu, G.: Detecting regime shifts: The causes of under- and overestimation. Management Science **51**(6) (2005) 932–947
14. Ghose, A., Smith, M.D., Telang, R.: Internet exchanges for used books: An empirical analysis of product cannibalization and welfare impact. Information Systems Research **17**(1) (2006) 3–19
15. Pauwels, K., Hanssens, D.: Windows of Change in Mature Markets. In: European Marketing Academy Conf., Braga, Portugal (2002)
16. Ng, A., Russell, S.: Algorithms for inverse reinforcement learning. In: Proc. of the 17th Int'l Conf. on Machine Learning, Palo Alto (2000) 663–670
17. Pardoe, D., Stone, P.: Bidding for Customer Orders in TAC SCM: A Learning Approach. In: Workshop on Trading Agent Design and Analysis at AAMAS, New York (2004) 52–58
18. Benisch, M., Greenwald, A., Grypari, I., Lederman, R., Naroditskiy, V., Tschantz, M.: Botticelli: A supply chain management agent designed to optimize under uncertainty. ACM Trans. on Comp. Logic **4**(3) (2004) 29–37
19. Ketter, W., Kryzhnyaya, E., Damer, S., McMillen, C., Agovic, A., Collins, J., Gini, M.: MinneTAC sales strategies for supply chain TAC. In: Int'l Conf. on Autonomous Agents and Multi-Agent Systems, New York (2004) 1372–1373
20. Kiekintveld, C., Vorobeychik, Y., Wellman, M.P.: An Analysis of the 2004 Supply Chain Management Trading Agent Competition. In: IJCAI 2005 Workshop on Trading Agent Design and Analysis, Edinburgh, Scotland (2005) 61–70
21. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. J. of the Royal Stat. Soc. **Series B, 39**(1) (1977) 1–38
22. Ketter, W.: Dynamic Regime Identification and Prediction Based on Observed Behavior in Electronic Marketplaces. In: Proc. of the Twentieth National Conference on Artificial Intelligence, Pittsburgh (2005) 1646–1647
23. Ketter, W., Collins, J., Gini, M., Gupta, A., Schrater, P.: A Computational Approach to Predicting Economic Regimes in Automated Exchanges. In: Proc. of the Fifteenth Annual Workshop on Information Technologies and Systems, Las Vegas, Nevada, USA (2005) 147–152