

We present basic building blocks of an agent that can use observable market conditions to characterize the microeconomic conditions of the market and predict future market trends. The agent can use this information to make both tactical decisions such as pricing and strategic decisions such as product mix and production planning. We develop methods that can learn dominant market conditions, such as over-supply or scarcity, from historical data using computational methods to construct price density functions. We discuss how this knowledge can be used, 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.

Detecting and Forecasting Economic Regimes in Automated Exchanges

Wolfgang Ketter^a, John Collins^a, Maria Gini^a, Alok Gupta^b, and Paul Schrater^a

^aDepartment of Computer Science and Engineering

University of Minnesota, Minneapolis, MN 55455, USA

{ketter,jcollins,gini,schrater}@cs.umn.edu

^bDepartment of Information and Decision Sciences

University of Minnesota, Minneapolis, MN 55455, USA

agupta@csom.umn.edu

1. Introduction

Business organizations seeking advantage are increasingly looking to automated decision support systems. These systems have become increasingly sophisticated in recent years. Advanced decision support systems are evolving into software agents that can act rationally on behalf of their users in a variety of application areas. Examples include procurement [27,7], scheduling and resource management [12,5], and personal information management [2,21].

In this paper, we show how machine learning techniques can be used to support rational decision making in a sales environment. We are particularly interested in environments that are constrained by capacity and materials availability. We demonstrate our approach in the context of an autonomous agent that is designed to compete in the Trading Agent Competition for Supply Chain Management (TAC SCM) [4].

Our method characterizes market conditions by distinguishable statistical patterns. We call these patterns *economic 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 [24], such predictions are rarely done for a micro economic environment.

In addition to the supply-chain trading example, there are a number of other domains that could benefit from our approach. Examples include agents for automated trading in financial markets, such as the Penn-Lehman Automated Trading Project [13], auction-based contracting environments, such as MAGNET [6], and other auctions, such as auctions for IBM PCs [20] or PDA's on eBay [9].

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. Literature review

Predicting prices is an important part of the decision process of agents or human decision makers. Kephart et al. [14] explored several dynamic pricing algorithms for information goods, where shopbots look for the best price, and pricebots adapt their prices to attract business. Wellman [32] analyzed and developed metrics for price prediction algorithms in the TAC Classic game, similar to what we have done for TAC SCM.

Massey and Wu [22] 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.

Ghose et al. [10] 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 [26], 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 [23] 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.

Chajewska, Koller, and Ormoneit [3] describe a method for predicting the future decisions of an agent based on its past decisions. They learn the agent's utility function by observing its behavior. Their approach is based on the assumption that the agent is a rational decision maker. According to decision theory, rational decision making amounts to maximization of the expected utility [31]. In TAC SCM, we cannot apply these techniques because the behaviors of individual agents are not directly observable.

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 [25], or by estimating the relationship between offer price and order probability with a linear cumulative density function (CDF) [1], or by using a reverse CDF and factors such as quantity and due date [15]. The Jackaroo team [33] applied a game theoretic approach to set offer prices, using a variation of the Cournot game for modeling the product market. The SouthamptonSCM [11] team used fuzzy reasoning to set offer prices. Similar techniques have been used outside TAC SCM to predict offer prices in first price sealed bid reverse auctions for IBM PCs [20] or PDA's on eBay [9].

In [17] the authors demonstrate a method for predicting future customer demand in the TAC SCM game environment, and use the predicted future demand to inform agent behavior. Their approach is specific to the TAC SCM situation, since it depends on knowing the formula by which customer demand is computed. Note that customer demand is only one of the factors for characterizing the multi-dimensional regime parameter space.

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 [16] 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. Tactical and strategic decisions

We are primarily interested in competitive market environments that are constrained by resources and/or production capacity. In such an environment, a manager who wants to maximize the value of available resources should be concerned about both strategic and tactical decisions. The basic strategic decision is to allocate available resources (financial, capacity, inventory, etc.) over some time horizon in a way that is expected to return the maximum yield in the market. For example, in a market that has a strong seasonal variation, one might want to build up an inventory of finished goods during the off season, when demand is low and prices are weak, in order to prepare for an expected period of strong demand and good prices. For the purpose of this paper, tactical decisions are concerned with setting prices to maximize profits, within the parameters set by the strategic decisions. So, for instance, if the forecast sales volume for the current week is 100,000 units, we would want to find the highest sales price that would move that volume.

We believe that our technique of modeling the economic regimes in a market can be used to inform both the strategic and tactical decision processes. In Figure 1 we show a schematic way how to present this process. In our formulation, a regime is essentially a distribution of prices over sales volume. We use this regime definition to characterize the market. By modeling an approximation of the probability of sale given an asking price and combining with demand numbers, this leads directly to (nearly) optimal pricing decisions.

In order to use our technique to inform the strategic decision process, we need to be able to forecast regime shifts in the market. If our forecast shows an upcoming period of low demand

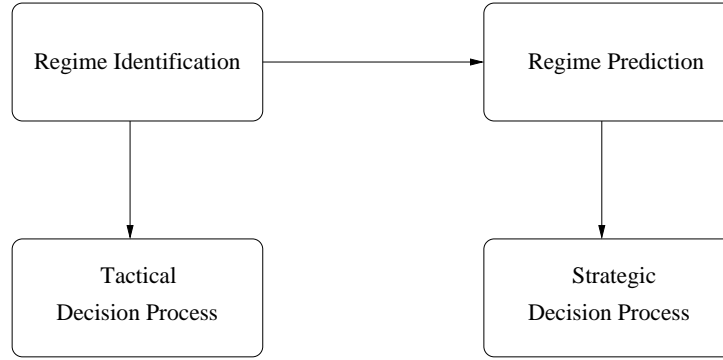


Figure 1. Process chart – Regime identification is a tool for tactical decision making and regime prediction is a tool for strategic decision making.

and weak prices, we may want to sell more aggressively in the short term, and we may want to limit procurement and production to prevent driving an oversupply into the market. On the other hand, if our forecast shows an upcoming period of high demand and strong prices, we may want to increase procurement and production, and raise short-term prices, in order to be well-positioned for the future.

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 oversupply. 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 oversupply 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 oversupply situations prices are lower. The agent should primarily control costs, and therefore either do pricing based on costs, or wait for better market conditions.

Figure 2 shows typical curves for the probability of receiving an order for a given offer price. The slope of the curve and its position changes over time. According to economic theory, high prices and a steep slope correspond to a situation of scarcity, where price elasticity is small,

while a less steep slope corresponds to a balanced market where the range of prices is larger.

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.

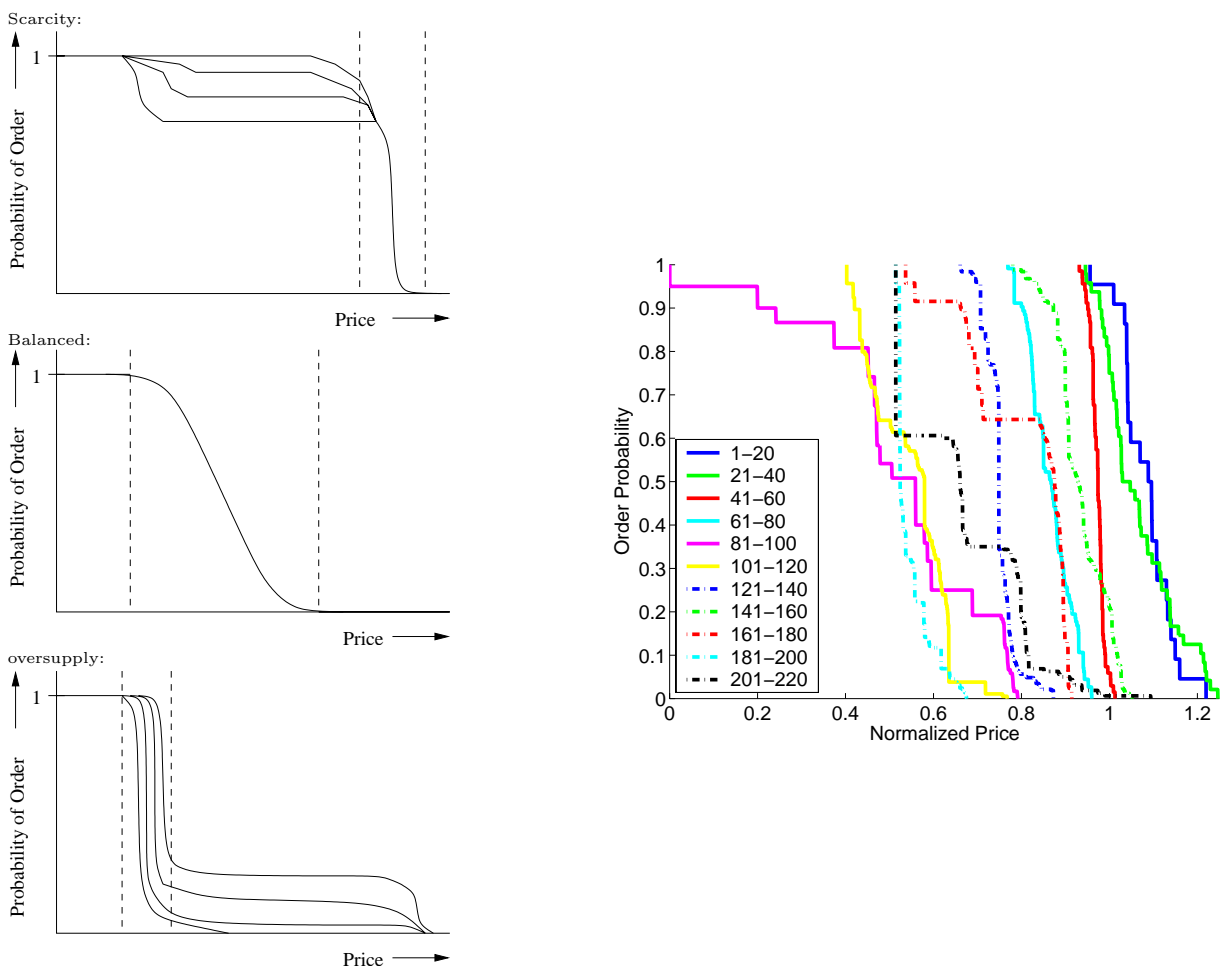


Figure 2. The reverse cumulative density function represents probability of order. Typical order probability curves during scarcity (left top), balanced (left middle) and oversupply (left bottom) regimes and experimental order probability curves (right). The curves shown on the right are from TAC SCM data and are measured at different days from the start of the game.

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:

$$np = \frac{ProductPrice}{NominalProductCost} \quad (1)$$

$$= \frac{ProductPrice}{AssemblyCost + \sum_{j=1}^{numParts} NominalPartCost_j} \quad (2)$$

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, $p(np)$, and to characterize regimes. We start by estimating the price density function by fitting a Gaussian mixture model (GMM) [30] 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.

In this paper, we present results using a GMM with fixed means, μ_i , and fixed variances, σ_i , since we want one set of Gaussians to work for all games off-line and online. We use the Expectation-Maximization (EM) Algorithm [8] to determine the prior probability, $P(c_i)$, of the Gaussians components of the GMM. The means, μ_i , are uniformly distributed and the variances, σ_i^2 , tile the space. Specifically variances were chosen so that adjacent Gaussians are two standard deviations apart.

The density of the normalized price can be written as:

$$p(np) = \sum_{i=1}^N p(np|c_i) P(c_i) \quad (3)$$

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

$$p(np|c_i) = p(np|\mu_i, \sigma_i) = \frac{1}{\sigma_i\sqrt{2\pi}} e^{\left[\frac{-(np-\mu_i)^2}{2\times\sigma_i^2}\right]} \tag{4}$$

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 3. While the choice of N , the number of Gaussians', in a GMM is arbitrary, the choice should reflect a balance between accuracy and computational overhead. By accuracy we mean predicted accuracy, which is not the same as fit accuracy. Creating a model with a very good fit to the observed data does not usually translate well into predictions. If the model has too many degrees of freedom the likelihood of overfitting the data is great. We chose $N = 16$ and $N = 25$ to show the effect of model flexibility on results for several prediction measures and illustrate their tradeoffs.




Figure 3. The price density density function, $p(np)$, (right y-axis) estimated by the Gaussian mixture model with 16 components 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:

$$P(c_i|\text{np}) = \frac{p(\text{np}|c_i) P(c_i)}{\sum_{i=1}^N p(\text{np}|c_i) P(c_i)} \quad \forall i = 1, \dots, N \quad (5)$$

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

$$\vec{\eta}(\text{np}) = [P(c_1|\text{np}), P(c_2|\text{np}), \dots, P(c_N|\text{np})]. \quad (6)$$

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

The intuitive idea of a regime as recurrent economic condition is captured by discovering price distributions that recur across days. We define *regimes* by clustering price distributions across days using the k-means algorithm with a similarity measure on both probability vectors $\vec{\eta}(\text{np}_j)$ and normalized prices np. The clusters found by this method correspond to frequently occurring price distributions with support on contiguous range of np.¹

The center of each cluster (ignoring the last component which contains the rescaled price information) is a probability vector that corresponds to regime $r = R_k$ for $k = 1, \dots, M$, where M is the number of regimes. Collecting these vectors into a matrix yields the conditional probability 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 4 we distinguish five regimes, which we can call extreme oversupply (R_1), oversupply (R_2), balanced (R_3), scarcity (R_4), and extreme scarcity (R_5). We decided to use five regimes instead of the three basic regimes which are suggested by economic theory because in this way we are able to isolate outlier regimes, such as extreme oversupply and extreme scarcity, in a market. Regimes R_1 and R_2 represent a situation where there is a glut in the market, i.e. an oversupply situation, which depresses prices. Regimes R_3 represents a balanced market situation, where

¹We have found that sometimes data points corresponding to specific regimes are close in probability space, but not in price space. Specifically it can happen that one regime dominates the extreme low and the extreme high price range, with different regimes in between. This regime is more difficult to interpret in terms of market concepts like oversupply or scarcity. To circumvent this problem we perform clustering in an augmented space formed by appending a rescaled version of np to the probability vector. Specifically, the mean of np is subtracted and np is scaled so that its standard deviation matches the largest standard deviation of the probability vectors.




Figure 4. An example of learned regime probabilities , $P(R_k|np)$, over normalized price np , for a sample market in TAC SCM after training.

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, $p(np|c_i)$, 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 :

$$p(np|R_k) = \sum_{i=1}^N p(np|c_i) P(c_i|R_k). \quad (7)$$

The probability of regime R_k dependent on the normalized price np can be computed using

Bayes rule as:

$$P(R_k|\text{np}) = \frac{p(\text{np}|R_k) P(R_k)}{\sum_{k=1}^M p(\text{np}|R_k) P(R_k)} \quad \forall k = 1, \dots, M. \quad (8)$$

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 4 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.

We model regime prediction as a Markov process. The last step is the computation of the Markov transition matrix to be used by the agent for regime prediction. We construct the Markov transition matrix, $\mathbf{T}_{\text{predict}}(r_{t+1}|r_t)$ 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 $\overline{\text{np}}_t$. 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.

$$\underset{1 \leq k \leq M}{\text{argmax}} \vec{P}(R_k|\overline{\text{np}}_t).$$

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 current time step.
2. prediction of the posterior distribution of regimes n time steps into the future, done recursively.

If the data for the current time-step is unavailable then we need to include after the correction operation a one time-step prediction to the current time before we do a prediction of future regime. The agent can use this forecast of regime transitions to drive its strategic decision process.

5. A case study: TAC SCM

The Trading Agent Competition for Supply Chain Management [4] (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 limited-capacity assembly facility, and must pay for warehousing its inventory. In addition, 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.

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

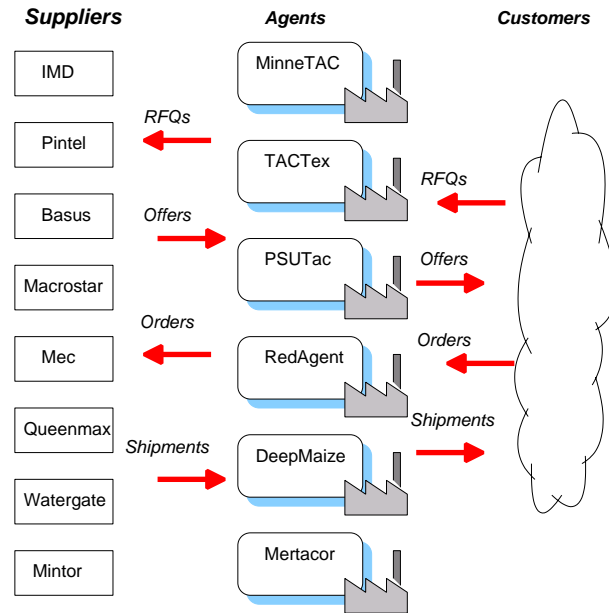


Figure 5. Schematic overview of a typical TAC SCM game scenario.

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 training² and 6 for testing³) 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. The mid-range price, \overline{np} , the price between the minimum and maximum, can be used to approximate the mean price, however, it does not always provide an accurate estimate of the mean price because of the local fluctuation in minimum and maximum prices. In other words, since the minimum and maximum prices could be unusual and temporary fluctuation, they may be outliers and not within the true distribution of the prevailing prices.

An example which shows how the mid-range value differs from the mean value is in Figure 6. 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 110, 120, 140 and at the end, 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

²3694@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.

³3717@tac3, 3718@tac3, 3724@tac3, 4253@tac4, 4254@tac4, 4260@tac4

that day's orders were sold at a much lower price.

To lower the impact of sudden price changes we implemented a Brown linear (i.e. double) exponential smoother with $\alpha = 0.5$. The general form of this smoother is:

$$S'_{t-1} = \alpha \cdot \overline{\text{np}}_{t-1} + (1 - \alpha) \cdot S'_{t-2} \quad (9)$$

$$S''_{t-1} = \alpha \cdot S'_{t-1} + (1 - \alpha) \cdot S''_{t-2} \quad (10)$$

$$\tilde{\text{np}}_{t-1} = 2 \cdot S'_{t-1} - S''_{t-1} \quad (11)$$

Since we only have the minimum and maximum prices from the previous day available and not the real mean, we decided to model $\tilde{\text{np}}_{t-1}$ as follows:

$$\tilde{\text{np}}_{t-1} = \frac{\tilde{\text{np}}_{t-1}^{\text{min}} + \tilde{\text{np}}_{t-1}^{\text{max}}}{2} \quad (12)$$

This results in a better approximation of the real mean price than smoothing only the mid-range price from the previous day. Figure 6 shows that the smoothed mid-range price, $\tilde{\text{np}}$, is closer to the mean price.

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

The mid-range price can be used to identify the corresponding regime online, as shown in Figure 7 (left). The data are from game 3721@tac3, which was not in the training set of games used to develop the regime definitions. The middle and right parts of Figure 7 show respectively the probability of receiving an order in a balanced and in a scarcity situation for different prices. Scarcity typically occurs early in the game and at other times when supply is low. These probabilities are computed from past game data for each regime.

Figure 8 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 could use this information to decide which strategy, or mixture of strategies, to follow.




Figure 6. Minimum, maximum, mean, mid-range, and smoothed mid-range daily normalized prices of computers sold, as reported during the game every day for the medium market segment in the 3721@tac3, one of the final games. The mean price is computed after the game using the game data, which include complete information on all the transactions.

A measure of the confidence in the regime identification is the entropy of the set S of probabilities of the regimes given the normalized mid-range price from the daily price reports \tilde{np}_{day} , where

$$S = \{P(R_1|\tilde{np}_{day}), \dots, P(R_M|\tilde{np}_{day})\} \quad (13)$$

and

$$\text{Entropy}(S) \equiv \sum_{k=1}^M -P(R_k|\tilde{np}_{day}) \log_2 P(R_k|\tilde{np}_{day}) \quad (14)$$

An entropy value close to zero corresponds to a high confidence in the current regime and an entropy value close to its maximum, i.e. for M regimes $\log_2 M$, indicates that the current market situation is a mixture of M almost equally likely regimes. An example for the medium market segment in game 3721@tac3 is shown in Figure 9.

Figure 10 (left) shows the factory utilization (FU) in percent, the ratio of offer to demand,

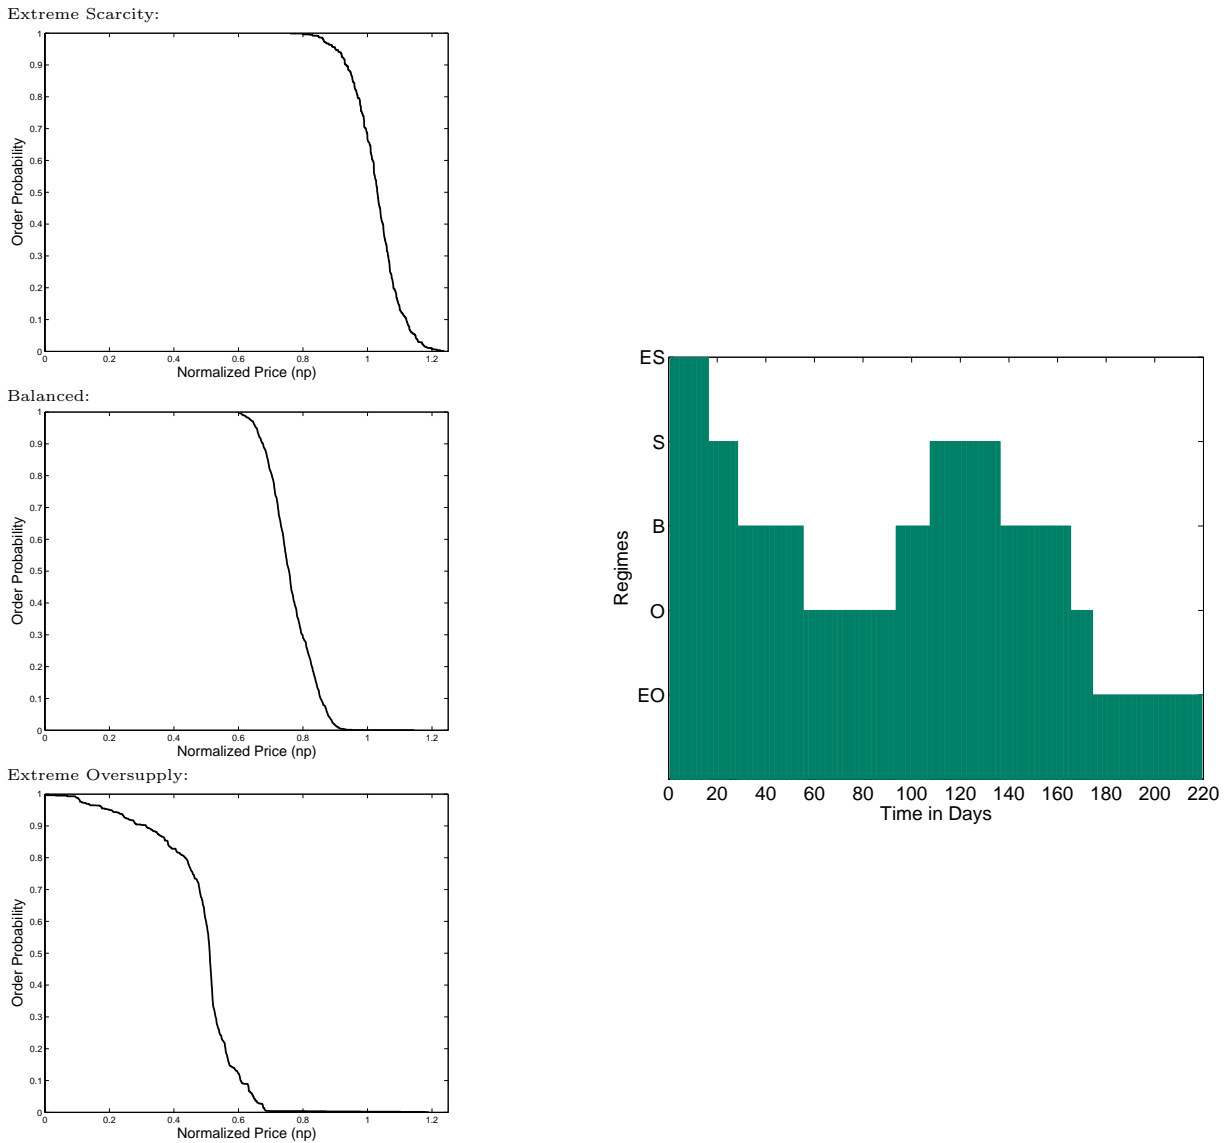


Figure 7. Game 3721@tac3 (Final TAC SCM 2005) – Regimes over time for the medium market computed online every day (left), probability of receiving an order by normalized price for an extreme scarcity situation (R_5 indicated by ES) (second from the left), for a balanced situation (R_3 indicated by B) (second from the right) and for an extreme oversupply situation (R_1 indicated by EO) (right).

which represents the proportion of the market demand that is satisfied, and the normalized price (np) over time. On the right side we display the quantity of the unsold finished goods inventory (FG) instead of factory utilization⁴. The regimes identified by our approach are superimposed,

⁴The quantity of the finished goods inventory is affected by other factors, such as storage cost, which have changed in the TAC SCM 2005 games. In 2005 and 2006 games agents tend to build to order and keep most of their inventory in the form of parts, not finished products.




Figure 8. Regime probabilities over time computed online every day for the medium market segment in game 3721@tac3 (Final TAC SCM 2005).




Figure 9. Daily entropy values of the five regimes for the medium market segment in game 3721@tac3. Notice how the entropy values match the regime probabilities shown in Figure 8.

where *ES* (or R_5) represents extreme scarcity, *S* (or R_4) scarcity, *B* (or R_3) balanced, *O* (or R_2) oversupply, and *EO* (or R_1) extreme oversupply. These factors clearly correlate with market




Figure 10. Game 3721@tac3 (Final TAC SCM 2005) – Relationships between regimes and normalized prices in the medium market. On the left axis, we show in the left figure the daily factory utilization and in the right figure the available finished goods inventory of all agents. In both figures we display on the left axis as well the ratio of offer to demand (which ranges from 0 to 5.38), which is scaled to fit between the minimum and maximum values of the left axis. On the right axis we show the normalized prices. The dominant regimes are labeled along the bottom.

regimes, but they are not directly visible to the agent during the game. For example, the figure shows that when the offer to demand ratio is high (i.e. oversupply) prices are low and vice versa. We can observe that the ratio of offer to demand changes significantly during the game. For instance, on day 111 the ratio of offer to demand is 1.95 and prices are high. On day 208 the ratio of offer to demand is much higher, 5.38, and prices are lower. We can also observe that prices tend to lag changes in ratio of offer to demand.

5.3. Predicting regime transitions

For TAC SCM, we model the prediction of future regimes as a Markov process. We construct a Markov transition matrix, $\mathbf{T}_{\text{predict}}(r_{t+1}|r_t)$ off-line by a counting process over past games. This matrix represents the posterior probability of transitioning in day $t+1$ to regime r_{t+1} given the current regime in day t , r_t .

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 smoothed mid-range normalized price \tilde{np} obtained since the time of the last regime change, t_0 , until the previous day, $t-1$. We

use $\vec{P}(r_{t-1}|\{\tilde{n}p_{t_0}, \dots, \tilde{n}p_{t-1}\})$, to indicate a vector of the posterior probabilities of all the regimes on day $t - 1$.

2. a prediction of regime posterior probabilities for the current day, t . The prediction of the posterior distribution of regimes n days into the future, $\vec{P}(r_{t+n}|\{\tilde{n}p_{t_0}, \dots, \tilde{n}p_{t-1}\})$, is done recursively as follows:

$$\begin{aligned} & \vec{P}(r_{t+n}|\{\tilde{n}p_{t_0}, \dots, \tilde{n}p_{t-1}\}) \\ &= \sum_{r_{t+n}} \cdots \sum_{r_{t-1}} \left\{ \vec{P}(r_{t-1}|\{\tilde{n}p_{t_0}, \dots, \tilde{n}p_{t-1}\}) \cdot \prod_{j=0}^n \mathbf{T}_{\text{predict}}(r_{t+j}|r_{t+j-1}) \right\} \end{aligned} \quad (15)$$

Examples of regime predictions for game 3721@tac3 for the medium market segment are shown in Figure 11 and Figure 12. 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 11. Regime predictions for game 3721@tac3 starting on day 80 for 20 days into the future for the medium 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.




Figure 12. Regime predictions for game 3721@tac3 starting on day 110 for 30 days into the future for the medium market segment.

Figure 11 shows a predicted change from from an oversupply situation to a balanced situation.

This means that the agent should sell less today and build up more inventory for future days when prices will be higher. On the other hand we see in Figure 12 a predicted change from the scarcity to the balanced regime. In this case the agent should try to sell more aggressively the current day, since prices will be decreasing in the next days.

6. Performance of regime identification and prediction

Our method is useful to the extent that it characterizes and predicts real qualities of the market. There are many hidden variables in a competitive market, such as the inventory positions and procurement arrangements of the competitors. Our method uses observable historical and current data to guide tactical and strategic decision processes. In this section we evaluate the practical value of regime identification and prediction.

6.1. Relationship between identified regime and market variables

We expect identified regimes to qualitatively represent the status of the important hidden market factors. A correlation analysis of market parameters of the training set is shown in Figure 13. The p-values for the correlation analysis are all less than 0.01. Regime *EO* (extreme oversupply) correlates positively with quantity of finished goods inventory, negatively with percent of factory utilization, positively with the ratio of offer to demand, and negatively with normalized price. On the other hand, in Regime *ES* (extreme scarcity) we observe a negative correlation with the amount of unsold finished goods inventory, with the percent of factory utilization, and the ratio of offer to demand, and positively with normalized price.

An advantage of using 5 regimes instead of 3 regimes is that we gain two degrees of freedom for better decision making, by isolating the outliers in the market. For example, regime *EO* (extreme oversupply) is different from Regime *O* (oversupply) since it presents a potential price war situation. Another difference between regime *EO* and regime *O* is that regime *EO* is universally unprofitable and that regime *O* is marginally profitable for most agents. Regimes *B* and *S* are universally profitable and in regime *ES* some agents have left the market. The major difference between the scarcity regime, *S*, and the extreme scarcity regime, *ES*, is that in regime *S* the factory runs at full capacity, caused by excess demand, and in regime *ES* we observe a scarcity of parts, with the result that production capacity is underutilized.




Figure 13. Training set (18 games) – Correlation coefficients between regimes and quantity of finished goods inventory, factory utilization, the ratio of offer to demand, and normalized price (np) in the medium market segment. All values are significant at the $p = 0.01$ level.

Another way to evaluate the quality of regime identification is given by an interpretation of the k-means clustering algorithm. Essentially, it finds points along the path that connects the regime centers in the regime probability space. In Figure 14 we represent the results of the k-means clustering algorithm, or the learned regime probabilities. For ease of visualization we use only 3 regimes to explain the learned behavior; the 5 regime case produces similar results, but they are harder to visualize. We can see that the learned regime probabilities in the posterior probability space connect the regimes in the “expected” way. In other words, we do not see points directly between scarcity and oversupply; instead, the path leads from scarcity through balance to oversupply.

Since the behavior of an agent should depend not just on the current regime but also on expected future regimes, the agent needs to predict future regimes. We would expect a dynamic regime prediction algorithm to move along this path of learned regime probabilities.

6.2. Prediction of discrete regime change

We measure the accuracy of regime prediction using a count of how many times the regime predicted is the correct one. As ground truth we measure regime switches and their time off-




Figure 14. An example of learned regime probabilities, $P(R_k|np)$, for the medium market segment in TAC SCM after training.

line using data from the game. Starting with day 1 until day 199, we forecast every day the regimes for the next 20 days and we forecast when a regime transition would occur. The reason for limiting the prediction to 20 days is that every 20 days the agent receives a report which includes the mean price of each of the computer types sold since the last market report, and so it can correct, if needed, its current regime identification. Experimental results for a GMM with 16 and 25 components are shown in Figure 15.




Figure 15. Success rate of correct regime shift prediction. The left figure is generated using a GMM with 16 components and the right figure with 25 components.

Table 1 reports the average number of regime changes and standard deviation for each market segment of the testing set. We see that the method produces robust results when varying the number of Gaussians in the GMM.

	low market avg/stdev	medium market avg/stdev	high market avg/stdev	# Gaussians
# regime changes	11.20/3.16	14.3/3.47	13.4/3.58	16
	11.25/3.56	13.5/5.2	12.75/3.44	25

Table 1

Average number of regime changes and standard deviation of the testing set.

6.3. Prediction of regime distribution

The above results are based on discrete regimes, i.e., using only the dominant regime of each predicted day to the actual value of any given day. One measure which can be used in determining the closeness of all individual predicted regime probabilities to the actual ones is called the *Kullback-Leibler (KL) divergence* [19,18]. This is a quantity which measures the difference between two probability distributions in bits, meaning the smaller the measure the closer the predictions are to optimal. We can calculate the Kullback-Leibler divergence, $KL(P(R)^{pred}||P(R)^{actual})$ as:

$$KL(P(R)^{pred}||P(R)^{actual}) = \sum_{r \in \mathcal{R}} P(R)^{pred} \log \left(\frac{P(R)^{pred}}{P(R)^{actual}} \right) \quad (16)$$

The KL difference can be interpreted in terms of how much additional data is needed to achieve optimal prediction performance. The precision of this data is given by the number of bits in the KL-divergence measure. For example a 1 bit difference would require an additional binary piece of information [28], like: “Were yesterday’s bids all satisfied?” If the difference between the two distributions is 0 then the predictions are optimal in sense that all the probabilistic information about pricing behavior is accurate (e.g. the predicted and actual distributions match).

In Figure 16 we show prediction results in terms of KL-divergence for a GMM with 16 components (left) and for a GMM with 25 components (right). Our predictions differ between 0.3 bits and 1 bits of information, as opposed to the Exponential Smoother predictions which vary

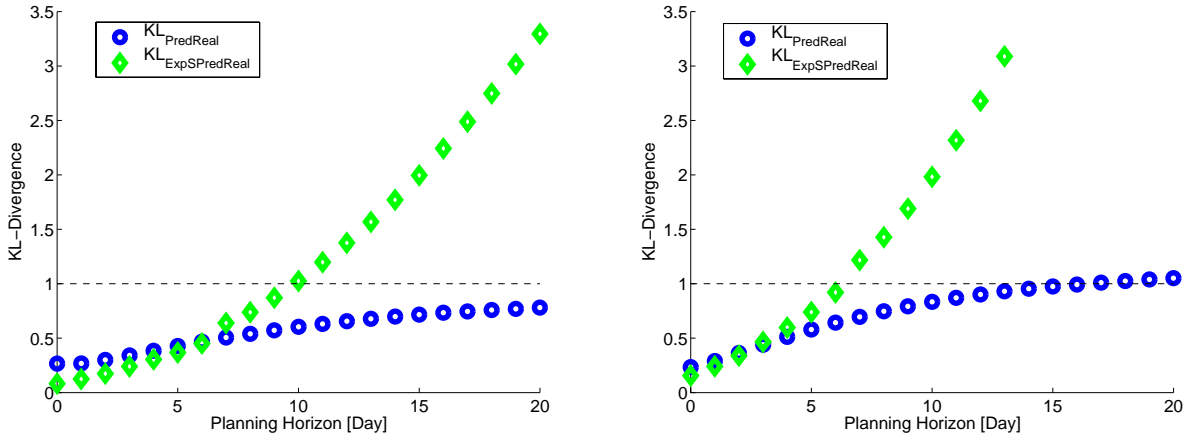


Figure 16. Normalized KL-divergence between the Markov predicted regime distribution and the actual distribution (circle), and the double exponentially smoothed predicted distribution and the actual distribution (diamond) over the planning horizon. KL-divergences computed using 5 regimes for the medium market segment over the testing set. The left figure is generated using a GMM with 16 components and the right figure with 25 components.

between 0.3 and 3.5 bits for a GMM with 16 components and between 0.2 and 6.5 (not shown to maintain the same scale for the KL-divergence with the right figure) bits for a GMM with 25 components. The 20 days, Exponential Smoother predictions are approximately 5.6 and 45 times as bad as Markov predictions. It is typically acceptable having a KL-divergence less than or close to one. There will not be significant gains by obtaining more information in the estimation procedure.

We observe that a the GMM with 25 components fits the actual regime probabilities much better for the first three days (tactical decision making) and the GMM with 16 components fits the actual regime probabilities much better in the long term (strategic decision making). The best estimate for the current day is given by the exponential smoother and as a result should be used as an input to generate the price density (explained in Section 7) and sales offer prices for the current day.

The KL-divergence tells us that the predicted regime distribution in the long term is closer to the real distribution with a GMM with 16 components than the one with 25 components. On the other hand we observe in Figure 15 the opposite is true. Here actually the GMM with 25 components has a higher regime change prediction accuracy. The reason is that we need to

optimize the number of GMM components for the right measure. The regime change success rate is based on a discrete regime identification, but the KL-divergence measures the closeness of the whole distribution. To match closely the whole distribution is much more important for automated applications, e.g. dynamic pricing algorithms, that utilizes the continuous distribution. A human decision maker on the other side might get a fast and deep insight from prediction of discrete future regimes, since they translate possible directly into some actions on the procurement or sales side.

6.4. Prediction error

We also measured the prediction error for the price distributions generated by our model. Because the mixture of Gaussians can only approximate the true price distributions, we measured the difference between observed price frequencies and model predictions using a Monte Carlo method. In particular, price frequencies were computed for 64 bins from game data to form an empirical histogram. Simulated price data was sampled from the mixture model and binned as per real data. Prediction error was defined as the 1-norm (sum of absolute differences) distance between simulated and measured histograms, averaged across 1000 simulated data samples. In Figure 17 we explain the algorithm used to analyze price predictions when sampling from the learned GMM with 16 and 25 components and table 2 displays the results for the fitted GMMs.

```

1 Inputs:
2    $pn_{avg}$ : original normalized price density
3    $numBins$ : the number of histogram bins
4    $GMM$ : learned Gaussian Mixture Model
5    $maxIter$ : number of iterations
6 Output:
7   PredErr: the overall mean prediction error
8 Process:
9   for  $j$  until  $maxIter$ 
10     $pn_{samp} = Monte\_Carlo\_Sampling(GMM)$ 
11     $Error(j) = \frac{|pn_{avg} - pn_{samp}|}{numBins}$ 
12  end loop
13   $meanErr = \overline{Error}$ 
14   $PredErr = numBins \cdot \sum meanErr$ 
15  return PredErr

```

Figure 17. Prediction error algorithm.

	low market	medium market	high market	# Gaussians
Prediction Error in %	6.69	6.89	8.99	16
Prediction Error in %	5.75	5.48	5.95	25

Table 2

Overall prediction error for a 16 and a 25 GMM in the three market segments. Results were obtained after averaging over 1000 iterations.

The results show that the total error introduced by the mixture model approximation varied between 5% – 8%, with more components resulting in slightly lower errors.

7. Prediction of price density and price trend

In our approach, an agent predicts the price distribution, see Equation 17, using the predicted regime distribution and the learned GMM for every day over the planning horizon n with a range of values for np .

$$\begin{aligned}
 & p(\text{Price at } t + n \mid \text{Price History since Regime Change}) \\
 & = f(\text{Predicted Regime Probabilities}, \text{Estimated Price Density})
 \end{aligned} \tag{17}$$

Figure 18 shows the forecast price density for game 3721@tac3, for 30 days starting at day 110. The dashed curve represents the price density for the first forecasted day, the thick solid line shows the price density for the last forecasted day, and the thin solid curves show the forecast for the intermediate days. As expected the predicted price density broadens as we forecast further into the future, reflecting a decreasing certainty in the prediction.

We can also compare the actual price trends with our predictions. Figure 19 shows the real mean price trend along with forecast price trends based on the different predictors, e.g. the 5%, 10% and the 50% percentiles. All the curves in the figure represent a relative price trend – to better compare the different predictors which each other graphically, we subtracted from each forecasted value the first predicted value, so that they all start at zero.

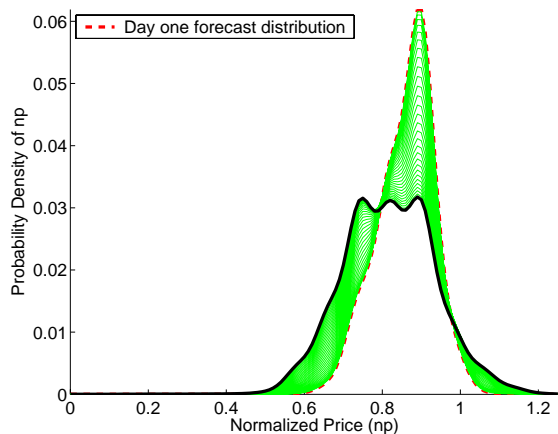


Figure 18. Predicted price density using the Markov model for game 3721@tac3 from day 110 until day 140 in the medium market segment. The red dashed curve is the price density estimate for the current day, the black solid curve is the price density estimate for the last day in the planning horizon, and the green solid curves are the estimates for the intermediate days.




Figure 19. Predicted price trend for game 3721@tac3 from day 110 until day 140 in the medium market segment. The solid curve is the real mean price and the dashed and dotted curves are predicted price trends based on the 5%, 10% and 50% percentile on the predicted price density.

8. 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.

8.1. Contributions

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 other agents. Our approach also learns the dynamics and durations of price regimes, and

when to expect a shift in the dominant regime. This is important information that is difficult to represent with regression-based methods. For example, regression in an expanding market (where prices increase) will extrapolate increasing prices using the slope of recent price data. On the other hand, the regime approach can learn that expansion (or scarcity) regimes are typically limited in duration and predictably followed by other regimes. When prices are increasing, it is more important to know if prices will fall by the end of the planning horizon, which can be invaluable information for a decision maker. Our method can 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.

8.2. Future directions

Our approach maintains the uncertainty in price prediction by maintaining a price distribution, which allows an agent to avoid over-committing to risky decisions. We intend to apply our method in other domains where predicting price distributions appears fruitful, including data from Amazon.com, eBay.com, and financial applications like stock tracking and forecasting.

We have implemented the regime identification and prediction method in a TAC SCM agent and integrated it into the overall decision making process. We are currently using regime predictions for strategic decision making in the sales component of our agent. We have begun to design an algorithm that use regime predictions for tactical decision making as well. Ultimately, we plan to combine probability information supplied by the algorithm with information about possible consequences of actions to optimize decision making. In particular, we would like to use the whole regime-based predicted price density to optimize tactical sales pricing by maximizing expected utility [31].

In addition, we plan to apply we plan to apply reinforcement learning [29] to map economic regimes to internal operational regimes and operational regimes to actions, such as procurement and production scheduling. Under operational regimes we understand a state which includes which actions to take next while knowing the current regime and receiving the regime forecast.

9. Appendix: Summary of notation

Symbol	Definition
np	Normalized price
\bar{np}	Mid-range normalized price
\tilde{np}^{min}	Smoothed minimum normalized price
\tilde{np}^{max}	Smoothed maximum normalized price
\tilde{np}	Smoothed mid-range normalized price
α	Smoothing coefficient
$p(np)$	Density of the normalized price
GMM	Gaussian Mixture Model
N	Number of Gaussians of the GMM
$p(np c_i)$	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
$P(c_i np)$	Posterior probability of the i -th Gaussian of the GMM given a normalized price np
$\vec{\eta}(np)$	N -dimensional vector with posterior probabilities, $P(c_i np)$, of the GMM
M	Number of regimes
R_k	k -th Regime, $k = 1, \dots, M$
$\mathbf{P}(c r)$	Conditional probability matrix (N rows and M columns) resulting from k-means clustering
$p(np R_k)$	Density of the normalized price np given a regime R_k
$P(R_k np)$	Probability of regime R_k given a normalized price np
t	Current time
t_0	Time of last regime change
$\mathbf{T}_{\text{predict}}(r_{t+1} r_t)$	Markov transition matrix

REFERENCES

1. Michael Benisch, Amy Greenwald, Ioanna Grypari, Roger Lederman, Victor Naroditskiy, and Michael Tschantz. Botticelli: A supply chain management agent designed to optimize under uncertainty. *ACM Trans. on Comp. Logic*, 4(3):29–37, 2004.
2. P. Berry, K. Conley, M. Gervasio, B. Peintner, T. Uribe, and N. Yorke-Smith. Deploying a personalized time management agent. In *Proc. of the Fifth Int’l Conf. on Autonomous Agents and Multi-Agent Systems*, Hakodate, Japan, May 2006.
3. Urszula Chajewska, Daphne Koller, and Dirk Ormoneit. Learning an agent’s utility function by observing behavior. In *Proc. of the 18th Int’l Conf. on Machine Learning*, pages 35–42, Lafayette, June 2001.
4. John Collins, Raghu Arunachalam, Norman Sadeh, Joakim Ericsson, Niclas Finne, and Sverker Janson. The Supply Chain Management Game for the 2005 Trading Agent Competition. Technical Report CMU-ISRI-04-139, Carnegie Mellon University, Pittsburgh, PA, December 2004.
5. John Collins, Corey Bilot, Maria Gini, and Bamshad Mobasher. Decision processes in agent-based automated contracting. *IEEE Internet Computing*, 5(2):61–72, March 2001.
6. John Collins, Wolfgang Ketter, and Maria Gini. A multi-agent negotiation testbed for contracting tasks with temporal and precedence constraints. *Int’l Journal of Electronic Commerce*, 7(1):35–57, 2002.
7. CombineNet. Sourcing solutions. http://www.combinenet.com/sourcing_solutions/, 2006.
8. A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. of the Royal Stat. Soc., Series B*, 39(1):1–38, 1977.
9. Rayid Ghani. Price prediction and insurance for online auctions. In *Int’l Conf. on Knowledge Discovery in Data Mining*, pages 411–418, Chicago, Illinois, August 2005.
10. Anindya Ghose, Michael D. Smith, and Rahul Telang. Internet exchanges for used books: An empirical analysis of product cannibalization and welfare impact. *Information Systems Research*, 17(1):3–19, 2006.
11. Minghua He, Alex Rogers, Esther David, and Nicholas R. Jennings. Designing and evaluating an adaptive trading agent for supply chain management applications. In *IJCAI 2005*

- Workshop on Trading Agent Design and Analysis*, pages 35–42, Edinburgh, Scotland, August 2005.
12. I2. Next-generation planning. http://i2.com/solution_library/ng_planning.cfm, 2006.
 13. Michael Kearns and Luis Ortiz. The Penn-Lehman Automated Trading Project. *IEEE Intelligent Systems*, pages 22–31, 2003.
 14. Jeffrey O. Kephart, James E. Hanson, and Amy R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, 2000.
 15. Wolfgang Ketter, Elena Kryzhnyaya, Steven Damer, Colin McMillen, Amrudin Agovic, John Collins, and Maria Gini. MinneTAC sales strategies for supply chain TAC. In *Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1372–1373, New York, July 2004.
 16. Christopher Kiekintveld, Yevgeniy Vorobeychik, and Michael P. Wellman. An Analysis of the 2004 Supply Chain Management Trading Agent Competition. In *IJCAI 2005 Workshop on Trading Agent Design and Analysis*, pages 61–70, Edinburgh, Scotland, August 2005.
 17. Christopher Kiekintveld, Michael P. Wellman, Satinder Singh, Joshua Estelle, Yevgeniy Vorobeychik, Vishal Soni, and Matthew Rudary. Distributed feedback control for decision making on supply chains. In *Int'l Conf. on Auto. Planning and Scheduling*, Whistler, BC, Canada, June 2004.
 18. Solomon Kullback. *Information Theory and Statistics*. Dover Publications, New York, 1959.
 19. Solomon Kullback and Richard A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.
 20. Richard Lawrence. A machine learning approach to optimal bid pricing. In *8th INFORMS Computing Society Conf. on Optimization and Computation in the Network Era*, Arizona, January 2003.
 21. Bill Mark and Raymond C. Perrault. Calo: Cognitive assistant that learns and organizes. <http://www.ai.sri.com/project/CALO>, 2006.
 22. Cade Massey and George Wu. Detecting regime shifts: The causes of under- and overestimation. *Management Science*, 51(6):932–947, 2005.
 23. Andrew Ng and Stuart Russell. Algorithms for inverse reinforcement learning. In *Proc. of the 17th Int'l Conf. on Machine Learning*, pages 663–670, Palo Alto, June 2000.
 24. Denise R. Osborn and Marianne Sensier. The prediction of business cycle phases: financial

- variables and international linkages. *National Institute Econ. Rev.*, 182(1):96–105, 2002.
25. David Pardoe and Peter Stone. Bidding for Customer Orders in TAC SCM: A Learning Approach. In *Workshop on Trading Agent Design and Analysis at AAMAS*, pages 52–58, New York, July 2004.
 26. Koen Pauwels and Dominique Hanssens. Windows of Change in Mature Markets. In *European Marketing Academy Conf.*, Braga, Portugal, May 2002.
 27. Tuomas Sandholm. Expressive commerce and its application to sourcing. In *Proc. of the Twenty-First National Conference on Artificial Intelligence*, pages 1736–1743, Boston, MA, July 2006. AAAI.
 28. Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423 and 623–656, July and October 1948.
 29. Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning - An Introduction*. The MIT Press, Cambridge, 1998.
 30. D. Titterton, A. Smith, and U. Makov. *Statistical Analysis of Finite Mixture Distributions*. Wiley, New York, 1985.
 31. John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, Princeton, N.J., 1947.
 32. Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, and Yevgeniy Vorobeychik. Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research*, 2003.
 33. Dongmo Zhang, Kanghua Zhao, Chia-Ming Liang, Gonelur Begum Huq, and Tze-Haw Huang. Strategic trading agents via market modeling. *SIGecom Exchanges*, 4(3):46–55, 2004.