

Detecting and Forecasting Economic Regimes in Automated Exchanges

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

*Dept. of Decision and Information Sciences, RSM Erasmus University, NL

Dept. of Computer Science and Engineering, University of Minnesota, USA

[†]Dept. of Information and Decision Sciences, University of Minnesota, USA

wketter@rsm.nl, {jcollins,gini,schrater}@cs.umn.edu, agupta@csom.umn.edu

Abstract

We show how an autonomous agent can use observable market conditions to characterize the microeconomic situation 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 to learn dominant market conditions, such as over-supply or scarcity, from historical data using Gaussian mixture models to construct price density functions. We discuss how this model can be combined with real-time observable information to identify the current dominant market condition and to forecast market changes over a planning horizon. We forecast market changes via both a Markov correction-prediction process and an exponential smoother. Empirical analysis shows that the exponential smoother yields more accurate predictions for the current and the next day (supporting tactical decisions), while the Markov correction-prediction process is better for longer term predictions (supporting strategic decisions). Our approach offers more flexibility than traditional regression based approaches, since it does not assume a fixed functional relationship between dependent and independent variables. We validate our methods by presenting experimental results in a case study, the Trading Agent Competition for Supply Chain Management.

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 [7], scheduling and resource management [15, 5], and personal information management [2, 25].

In this paper, we show how machine learning techniques can be used to support rational decision making in competitive sales environments. We are particularly interested in environments that are constrained by capacity and materials availability, and that are characterized by limited visibility of significant factors, such as the inventory positions of competitors. Such environments exist in virtually all manufacturing industries. They are more interesting for our present study than other environments where either capacity or materials may not be a constraint. For example, in the airline industry capacity is constrained but material is typically not a major constraint. In financial markets, neither capacity nor material are real constraints. We focus on the more complex case because it is easier to transfer results from more complex situations to less complex ones.

Our method characterizes market conditions by detecting distinguishable statistical patterns in historical data and using those patterns to classify the range of economic environments that a system

is expected to encounter in the future. By distinguishable pattern we mean a computationally identifiable separation in statistical properties such as price, inventory, or demand. These patterns are obtained by clustering historical data, and identifying, via a semi-supervised approach, clusters that reflect qualitatively distinct economic conditions. We call these patterns *economic regimes*.

We outline how to identify regimes from observable data and how to 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 [28], such predictions are rarely done for a micro economic environment.

Because economic regimes reflect qualitative distinctions in the economic environment, they can be used to drive decision criteria for both short-term (tactical) and medium-to-long term (strategic) decisions. For example, suppose we are in an economic environment where there is a balance in supply and demand, but we predict, using our method, that supply will outstrip demand in the near future. We might decide to lower our prices to reduce inventory and slow the acquisition of raw materials. However, if we predict that in future demand will outstrip supply, we may want to raise prices and increase our acquisition of raw material. Therefore, identification and prediction of regimes can help with tactical (pricing) and strategic (procurement) planning .

After a review of relevant literature, we describe the information needed to make strategic and tactical sales decisions. This is followed by a discussion of the concept of “economic regimes” and their representation using learned probability density functions. 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]. A TAC SCM agent must make resource allocation and pricing decisions using the limited visible 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.

While it is conceivable that in future many operations will be conducted with greater levels of autonomy by computer systems (similar to stock trading on Nasdaq), such systems are not extensively deployed in other business environments. TAC SCM affords a rich but risk free environment for testing economic decision processes. A simulated environment, such as TAC SCM, has also the advantage that we can set it up to enable repeating experiments with different algorithms and parameters [32]. This allows us to tease apart subtle interactions among decision processes in the agent and in its interactions with the other agents that would be hard to find and study in the real world.

In addition to the supply-chain trading example we present here, there are many 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 [16], auction-based contracting environments, such as MAGNET [6], and other auctions, such as auctions for IBM PCs [24] or PDA’s on eBay [11].

Our approach contributes to the research and development of Smart Business Networks (SBN) [36, 35] since it creates value through providing regimes as a tool to gather network knowledge and to facilitate decision making, both for humans and for autonomous software agents. Simulations like TAC SCM help to construct and test economic models of real world SBNs.

For reading convenience, we present a summary of our notation in the Appendix.

2 Literature review

Human and automated decision makers are often faced with the need to predict prices. Kephart et al. [17] explored several dynamic pricing algorithms for information goods, where shopbots look for the best price, and pricebots adapt their prices to attract business. Wellman [38] compared approaches used by agents to predict hotel and flight prices in the TAC Classic game. He analyzed the approaches and developed metrics, such as the Euclidean distance between the predicted and the actual price.

Massey and Wu [26] show in their analysis that the ability of decision makers to correctly identify the onset of a new regime can be 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, since they have difficulties in distinguishing the signal from the noise.

Ghose et al. [12] 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 [30], 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 Russell [27] 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 or indeed most real-world markets.

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 of the minimum and maximum daily prices [29], 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 [19]. The Jackaroo team [39] applied a game theoretic approach to set offer prices, using a variation of the Cournot game for modeling the product market. The SouthamptonSCM [13] 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 [24] or PDA's on eBay [11].

In [21] 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 also that customer demand is only one of the factors for characterizing the multi-dimensional regime parameter space, and is not by itself a good short-term predictor of prices.

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 a consistent structure. Our approach models variability directly and does not assume a well-defined functional relationship.

An analysis [20] of the TAC SCM 2004 competition shows that supply and demand (which are 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 or production capacity or both. In such environments, 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 how to allocate the available resources (financial, production capacity, inventory, etc.) over some time horizon in a way that is expected to return the maximum yield. 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 high demand and strong prices.

For the purpose of this paper, tactical decisions are concerned with setting prices to maximize profit, 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. According to efficient markets theory [10], prices subsume information about customer demand and other relevant factors. Therefore, with a good method for predicting prices we implicitly account for the other factors that are encoded in the price.

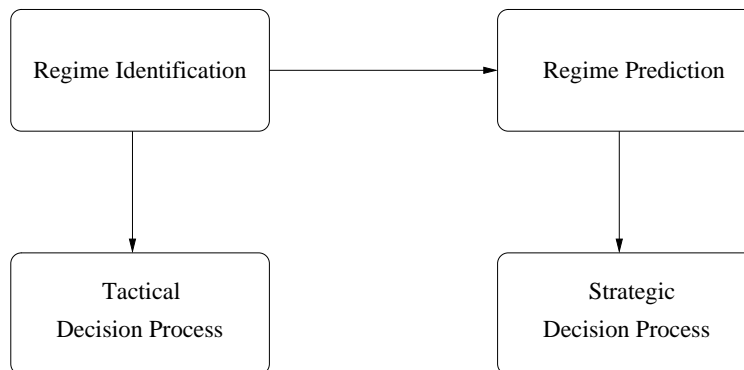


Figure 1: Process chart – Regime identification is a tool for tactical decision making. Regime prediction is a tool for strategic decision making.

Our technique of modeling economic regimes can be used to inform both the strategic and tactical decision processes. Figure 1 shows in a schematic way how on-line regime identification and prediction can support these processes.

In our formulation, a regime is essentially a distribution of prices over sales volume. at the tactical level, by modeling the probability of selling a product at a given price and combining that probability with demand numbers, we obtain directly (nearly) optimal pricing decisions. In other words, we are able to find the highest prices the market will support for a given sales volume.

In order to take advantage of regimes to inform the strategic decision process, we need to forecast regime shifts in the market. In our formulation, regime prediction is done either as a Markov prediction-correction process or using a double exponential smoother. as we will show later in Section 4.3. If the forecast shows an upcoming period of low demand 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 the market into oversupply. 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 in procurement, production planning, and product pricing. Different market conditions relate to different points on the supply and demand curves [14]. The law of supply and demand states that the equilibrium market price of a product is the intersection of consumer demand curve and producer supply curve.

We define a “scarcity” condition [8] when demand is elastic relative to supply, giving suppliers pricing power. This commonly occurs when market demand exceeds supplier capacity, at least temporarily. Oversupply occurs when supply is elastic relative to demand, and prices are not strongly affected by demand. Faced with an oversupply situation, the agent should primarily control costs, and therefore either price based on costs, or wait for better market conditions. A “balanced” market is not dominated by effects of elasticity. In balanced situations, prices are affected by variations in both demand and supply, so the agent has a range of options for maximizing expected profit.

Our approach based on regimes can be used to support decisions in both procurement and sales markets. In a procurement market, we may have little or no control over the availability of parts, but we can control the usable supply to a certain degree. When there is scarcity of parts, the prices for parts will increase. It is common that scarcity of parts results from excess demand for parts, which is likely to occur when demand for associated products is high. This is precisely why prediction of regimes is important. If we can predict future demand (via prices) well, then we may decide to acquire materials early, thereby reducing prices for material and increasing our profit margin.

Rapidly changing market conditions are often observed with commodities, such as oil or semiconductor chips. Such markets often experience high price volatility. For instance, the price of oil might go up due to shortage in supply or to inventories build up due to a softening demand in Asia. We believe that even though a market may be constantly changing, there are some underlying dominant patterns or economic regimes that characterize market conditions. Agents can guide their decision processes by mapping observed market conditions to a learned regime model and by predicting regime shifts over a planning horizon, as we shall describe in the following sections.

We use the term “order probability” with respect to price to represent the probability that an active customer (a customer who accepts an offer from some seller) will accept an offer at a particular price. If the market is in a scarcity condition, order probability remains high up to relatively high profit margins. The contrary is true if the market is in an oversupply situation. In this case there is a glut in the market (highly elastic supply), profit margins are low, and a small price increase will dramatically reduce the probability of order. In balanced situations the probability of order curve is between those two extremes. Figure 2 shows curves for the probability of receiving an order by offer price depending on the regime (left) or the time in the market (right). As we can see in the figure, the slope of the curve and its position changes over time as market conditions change.

Our approach includes (1) an off-line learning phase, where historical data are used to characterize the regimes, and (2) an online phase where current market data are used to identify the dominant regime and to make predictions.

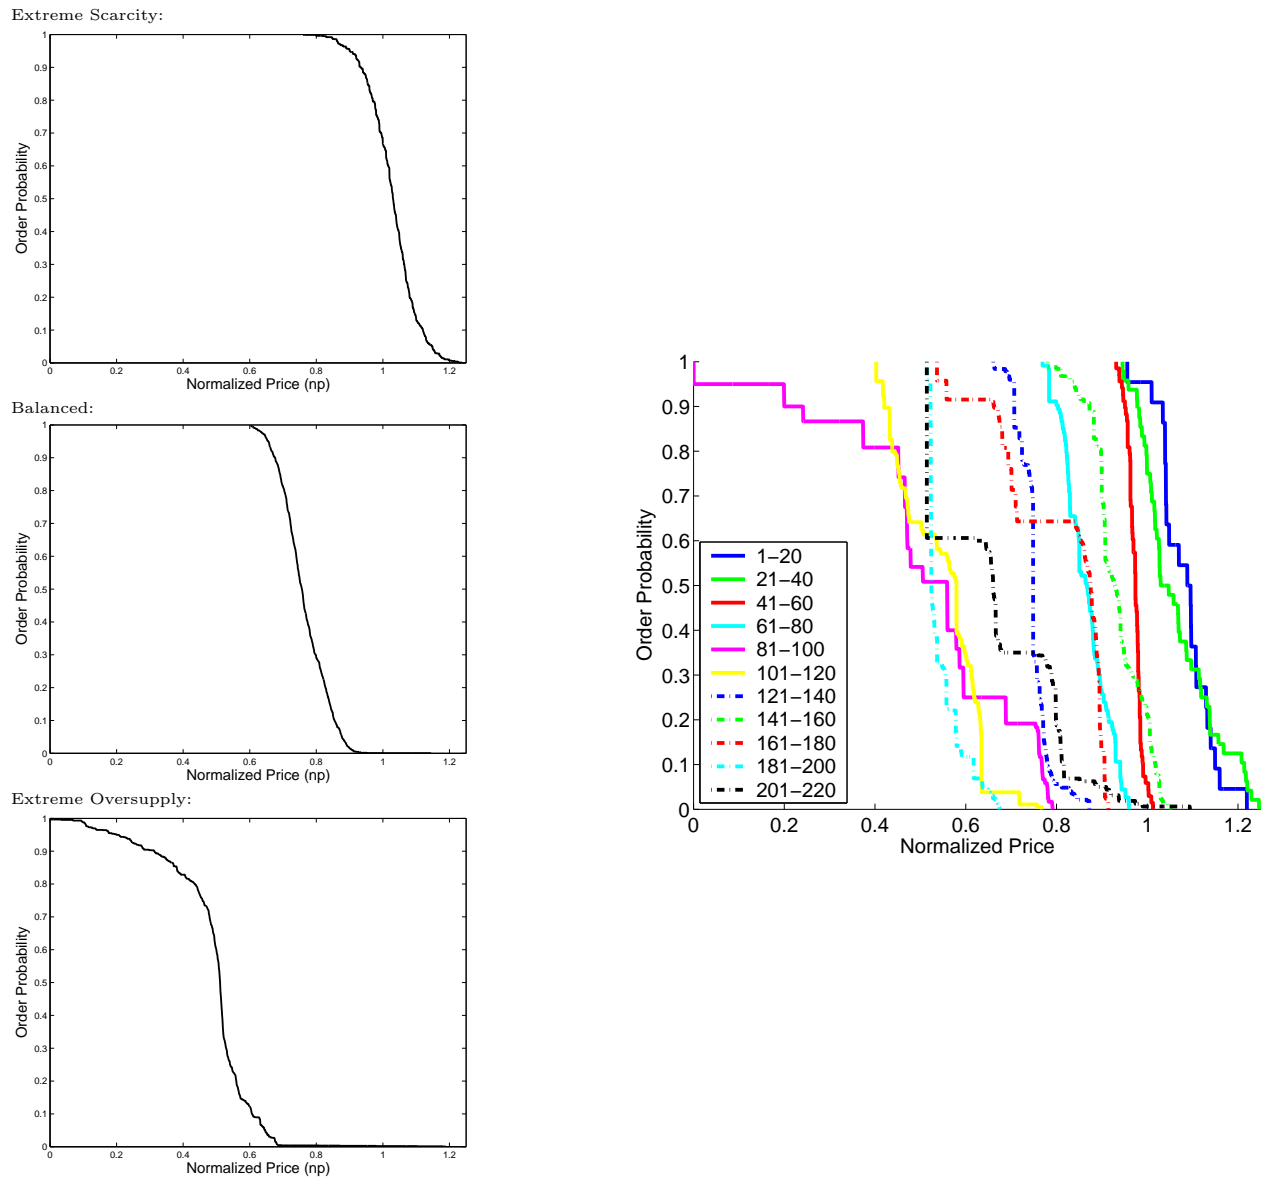


Figure 2: The reverse cumulative price density function represents the probability of order by price. The figure shows order probability curves for the medium market in a TAC SCM game during regimes of extreme scarcity (left top), balanced (left middle) and extreme oversupply (left bottom). The figure also shows (right) experimental order probability curves, computed every twenty days, during one of the games.

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.

Our model makes three specific assumptions about availability and quality of data. First, we assume that sufficient, relevant historical data are available for analysis. Second, we assume that the historical data are sufficiently representative of present and future market conditions. Third, to support online identification of regimes, we assume that some sort of price information is available

in real time.

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

$$\begin{aligned} \text{np}_g &= \frac{\text{ProductPrice}_g}{\text{NominalProductCost}_g} \\ &= \frac{\text{ProductPrice}_g}{\text{AssemblyCost}_g + \sum_{j=1}^{\text{numParts}} \text{NominalPartCost}_{g,j}} \end{aligned} \quad (1)$$

where $\text{NominalPartCost}_{g,j}$ is the nominal cost of the j -th part for good g , numParts is the number of parts needed to make the good g , and AssemblyCost_g is the cost of manufacturing the good g . An advantage of using normalized prices is that we can easily compare price patterns across different products. In the following, for simplicity of notation, we use np instead of np_g .

Historical data are used to estimate the price density, $p(\text{np})$, and to characterize regimes. We start by fitting a Gaussian mixture model (GMM) [34] to historical normalized price 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 cases off-line and online. We use the Expectation-Maximization (EM) Algorithm [9] to determine the prior probability, $P(\zeta_i)$, of the Gaussian 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:

$$p(\text{np}) = \sum_{i=1}^N p(\text{np}|\zeta_i) P(\zeta_i) \quad (2)$$

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

$$p(\text{np}|\zeta_i) = p(\text{np}|\mu_i, \sigma_i) = \frac{1}{\sigma_i \sqrt{2\pi}} e^{\left[\frac{-(\text{np} - \mu_i)^2}{2\sigma_i^2} \right]} \quad (3)$$

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 prediction accuracy, which is not the same as fit accuracy. Creating a model with a very good fit to the observed data does not always translate well into predictions. The more degrees of freedom the model has, the higher is the likelihood of overfitting the data.

Using Bayes' rule we determine the posterior probabilities for each Gaussian ζ_i :

$$P(\zeta_i|\text{np}) = \frac{p(\text{np}|\zeta_i) P(\zeta_i)}{\sum_{i=1}^N p(\text{np}|\zeta_i) P(\zeta_i)} \quad \forall i = 1, \dots, N \quad (4)$$

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

$$\vec{\eta}(\text{np}) = [P(\zeta_1|\text{np}), P(\zeta_2|\text{np}), \dots, P(\zeta_N|\text{np})]. \quad (5)$$




Figure 3: The price density function, $p(np)$, (right y-axis) estimated by a 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.

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

The intuitive idea of a regime as a recurrent economic condition is captured by discovering price distributions that recur across time periods in the market. We define *regimes* by clustering price distributions over time periods using the k -means algorithm with a similarity measure on the probability vectors $\vec{\eta}(np_j)$ and normalized prices np . The clusters found by this method correspond to frequently occurring price distributions with support on contiguous ranges 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 $\mathbf{P}(\zeta|r)$. The matrix has N rows, one for each component of the GMM, and M columns, one for each regime.

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

$$p(np|R_k) = \sum_{i=1}^N p(np|\zeta_i) P(\zeta_i|R_k). \quad (6)$$

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.

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

$$P(R_k|\text{np}) = \frac{p(\text{np}|R_k) P(R_k)}{\sum_{k=1}^M p(\text{np}|R_k) P(R_k)} \quad \forall k = 1, \dots, M. \quad (7)$$

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. Regimes provide a better assessment of market conditions.

In our experiments we tried different numbers of Gaussians for the GMM, typically 16 and 25, as discussed later in Section 6, and different numbers of regimes, typically 3 or 5. We found out that the number of regimes does not significantly affect the results regarding price trend predictions. We also tried k -means clustering with different initial conditions, but it consistently converged to the same results. More details are given in [18].




Figure 4: An example of learned regime probabilities , $P(R_k|\text{np})$, over normalized price np, with five regimes for a sample market in TAC SCM after training.

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 (price war) and extreme scarcity (temporary monopoly), 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 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.

4.2 Online identification of dominant regime

After historical sales data have been used to learn the characterization of different market regimes, an agent can use this information in real-time to identify the dominant regime. This can be done by calculating (or estimating) the normalized prices for the current time step. Time is discretized either in days, like in TAC SCM, or in other time units relevant to the domain.

Since complete current price information might not be available, we indicate the estimated normalized price at time t by $\widetilde{\text{np}}_t$. Depending on the application domain, the price estimate can be accurate, or can be an approximation, as shown later for TAC SCM in Section 5.2.

The current dominant regime is the one which has the highest probability, i.e.

$$\widetilde{R}_c \text{ s.t. } c = \underset{1 \leq k \leq M}{\operatorname{argmax}} \vec{P}(R_k | \widetilde{\text{np}}_t). \quad (8)$$

4.3 Online regime prediction

We model the prediction of future regimes in two different ways:

1. as a Markov correction-prediction process.
2. using an exponential smoother.

4.3.1 Markov correction-prediction

We construct a Markov transition matrix, $\mathbf{T}(r_{d+1}|r_d)$, which represents the posterior probability of transitioning at time $d + 1$ to regime r_{d+1} given the current regime r_d at time d . The matrix is computed off-line by a counting process over historical data.

The prediction is based on two distinct operations:

1. a correction (recursive Bayesian update) of the posterior probabilities for the regimes based on the value of the normalized price (if available, or of its estimate) obtained from the first measurement until the previous time, $d - 1$. We use $\vec{P}(r_{d-1} | \{\widetilde{\text{np}}_1, \dots, \widetilde{\text{np}}_{d-1}\})$, to indicate a vector of the posterior probabilities of all the regimes at time $d - 1$.
2. a prediction of regime posterior probabilities for the current time, d . The prediction of the posterior distribution of regimes n time units into the future, $\vec{P}(r_{d+n} | \{\widetilde{\text{np}}_1, \dots, \widetilde{\text{np}}_{d-1}\})$, is done by multiplying the prediction matrix n times, as follows:

$$\begin{aligned} & \vec{P}(r_{d+n} | \{\widetilde{\text{np}}_1, \dots, \widetilde{\text{np}}_{d-1}\}) \\ &= \sum_{r_{d+n}} \cdots \sum_{r_{d-1}} \left\{ \vec{P}(r_{d-1} | \{\widetilde{\text{np}}_1, \dots, \widetilde{\text{np}}_{d-1}\}) \cdot \prod_{j=0}^n \mathbf{T}_1^{\text{h}+1}(r_{d+j} | r_{d+j-1}) \right\} \end{aligned} \quad (9)$$

where

$$\mathbf{T}_1^{\text{h}+1}(r_d | r_{d-1}) = \prod_{n=0}^h \mathbf{T}_1(r_d | r_{d-1}) \quad (10)$$

4.3.2 Exponential smoother prediction

As an alternative to the Markov prediction process, we describe a method for regime predictions based on exponentially smoothed price predictions. We assume that price observations are only available for the past, up to the previous period $d - 1$. We first calculate the price trend, tr_{d-1} using a Brown linear exponential smoother [3]:

$$tr_{d-1} = \frac{\alpha}{1 - \alpha} \cdot (\widetilde{np}'_{d-1} - \widetilde{np}''_{d-1}) \quad (11)$$

where

$$\widetilde{np}'_{d-1} = \alpha \cdot \widetilde{np}_{d-1} + (1 - \alpha) \cdot \widetilde{np}'_{d-2} \quad (12)$$

$$\widetilde{np}''_{d-1} = \alpha \cdot \widetilde{np}'_{d-1} + (1 - \alpha) \cdot \widetilde{np}''_{d-2} \quad (13)$$

We can then estimate normalized prices for current and future periods as

$$\widetilde{np}_{d+n} = \widetilde{np}_{d-1} + (1 + n) \cdot \widetilde{tr}_{d-1}, \quad \forall n = 0, \dots, h \quad (14)$$

Given the current and future price estimates, we can use Equation 7 to predict future regimes.

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 (see Figure 5). An agent in TAC SCM is a software entity that represents a virtual manufacturer. During a game, agents make fully autonomous decisions without any human intervention. Over fifty academic or industrial research groups world wide have implemented agents. A game 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. Each agent begins with an empty bank account. The agent with the highest bank balance at the end of the game wins.

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 vary between games.

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.

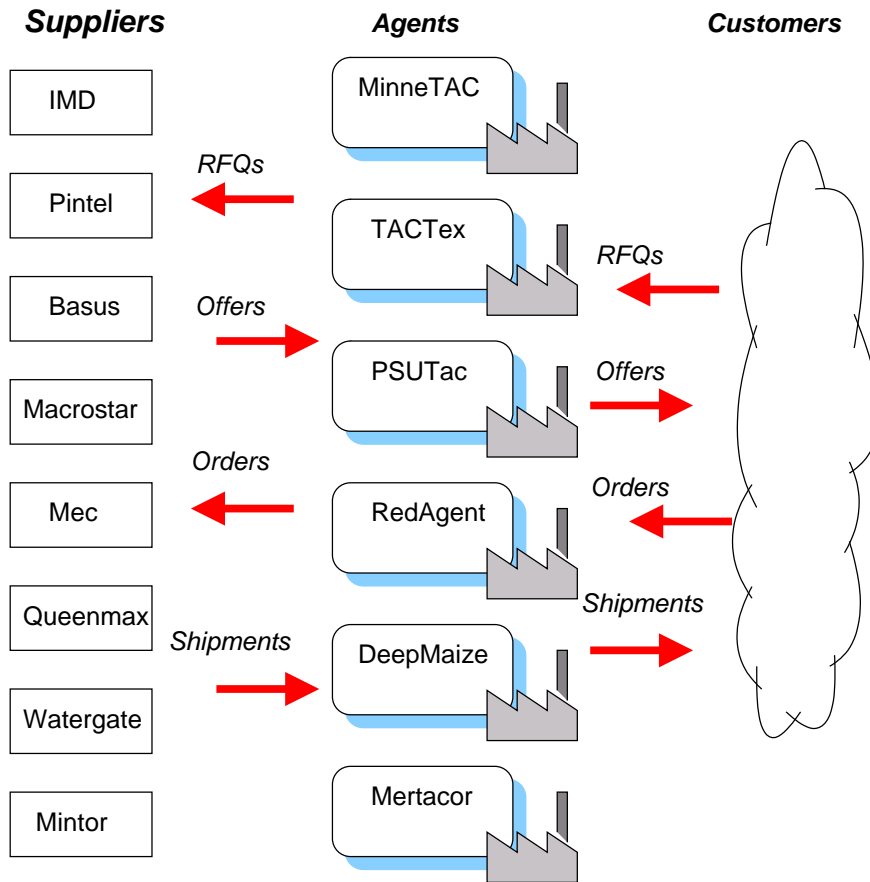


Figure 5: Schematic overview of a typical TAC SCM game scenario.

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.

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, as shown in Equation 1.

5.2 Estimate of prices of previous day sales

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 mean of 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 local fluctuations 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.




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.

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

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

To reduce the impact of fluctuations in minimum and maximum prices, we compute the smoothed mid-range price $\widetilde{\text{np}}_d$ on day d as the average of the smoothed minimum normalized price $\widetilde{\text{np}}_d^{\text{min}}$ and the smoothed maximum normalized price $\widetilde{\text{np}}_d^{\text{max}}$ for the same day:

$$\widetilde{\text{np}}_d = \frac{\widetilde{\text{np}}_d^{\text{min}} + \widetilde{\text{np}}_d^{\text{max}}}{2} \quad (15)$$

To lower the impact of sudden price changes we smooth the minimum normalized price (and, in a similar way, the maximum normalized price) using a Brown linear (i.e. double) exponential smoother ([3]) with $\alpha = 0.5$:

$$\widetilde{\text{np}}_d^{\text{min}} = 2 \cdot \widetilde{\text{np}}_d^{\text{min}'} - \widetilde{\text{np}}_d^{\text{min}''} \quad (16)$$

where

$$\widetilde{\text{np}}_d^{\text{min}'} = \alpha \cdot \text{np}_d^{\text{min}} + (1 - \alpha) \cdot \widetilde{\text{np}}_{d-1}^{\text{min}'} \quad (17)$$

$$\widetilde{\text{np}}_d^{\text{min}''} = \alpha \cdot \widetilde{\text{np}}_d^{\text{min}'} + (1 - \alpha) \cdot \widetilde{\text{np}}_{d-1}^{\text{min}''} \quad (18)$$

This results in a better approximation of the mean price than smoothing only the mid-range price from the previous day, as shown the example in Figure 6.

5.3 Online identification of dominant regime

During the game, the agent estimates on day d the current regime by calculating the smoothed mid-range normalized price $\widetilde{\text{np}}_{d-1}$ 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. $\text{argmax}_{1 \leq k \leq M} \vec{P}(R_k | \widetilde{\text{np}}_{d-1})$.

Figure 7 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.

A measure of the confidence in the regime identification is the entropy of the set S of the probabilities of the regimes obtained from the daily price reports $\widetilde{\text{np}}_d$, where

$$S = \{P(R_1 | \widetilde{\text{np}}_d), \dots, P(R_M | \widetilde{\text{np}}_d)\} \quad (19)$$

and

$$\text{Entropy}(S) \equiv \sum_{k=1}^M -P(R_k | \widetilde{\text{np}}_d) \log_2 P(R_k | \widetilde{\text{np}}_d) \quad (20)$$

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




Figure 7: Regime probabilities over time computed online every day for the medium market segment in game 3721@tac3 (Final TAC SCM 2005).




Figure 8: 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 7.

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 8. Since a market is usually dominated by two regimes, i.e. scarcity and balanced, an alternative measure is the ratio of the highest regime probability divided by the second highest regime probability. A more sophisticated approach to this measure can be found in [18].

5.4 Relationship between regimes and market variables

For our experiments we selected the number of regimes, after examining the data and looking at economic analyses of market situations. The regimes we chose are: *EO* (or R_1) which represents extreme oversupply, *O* (or R_2) oversupply, *B* (or R_3) balanced, *S* (or R_4) scarcity, and *ES* (or R_5) extreme scarcity,

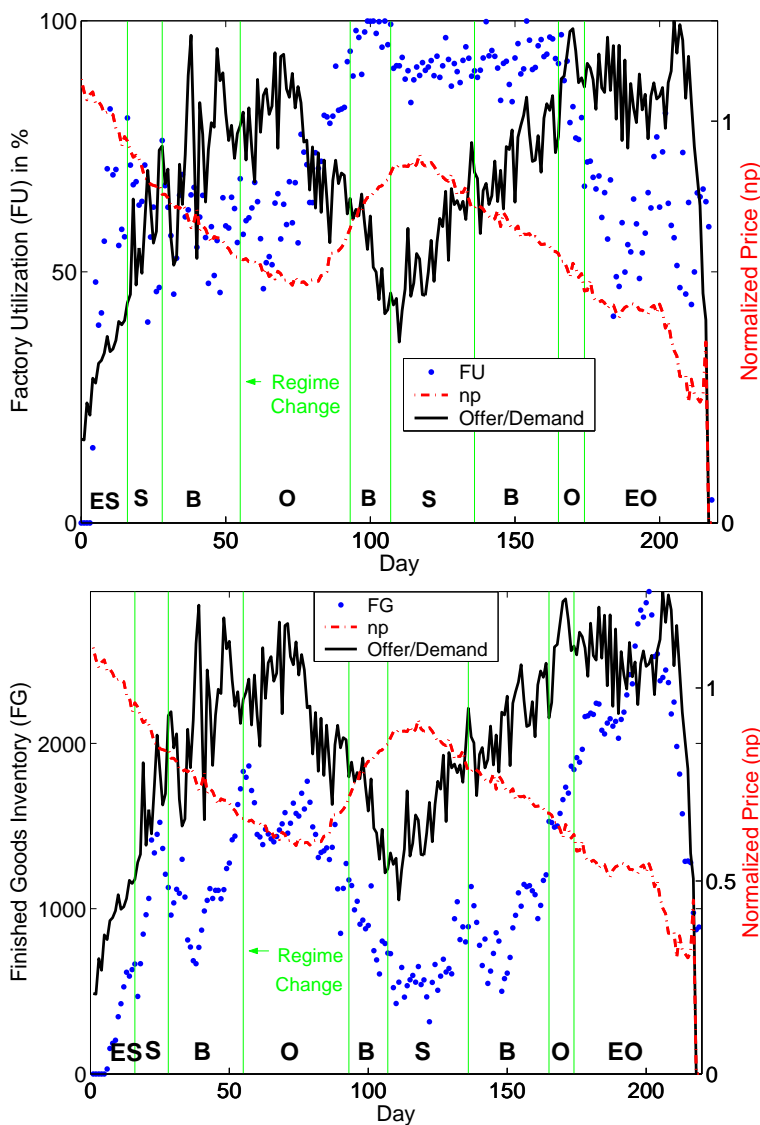


Figure 9: Game 3721@tac3 (Final TAC SCM 2005) – Relationships between regimes and normalized prices in the medium market. On the left axis, we show in the top figure the daily factory utilization (FU) and in the bottom figure the available finished goods inventory (FG) of all the agents. In both figures we show the ratio of offer to demand (which ranges from 0 to 5.38), scaled to fit between the minimum and maximum values on the left axis. On the right axis we show the normalized prices. The dominant regimes are labeled along the bottom.

Figure 9 (top) shows the factory utilization (FU) in percent, the ratio of offer to demand, which

represents the proportion of the market demand that is satisfied, and the normalized price over time. On the bottom figure we display the quantity of the unsold finished goods inventory (FG) instead of factory utilization⁴.

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. These factors clearly correlate with market regimes, but they are not directly visible to the agent during the game.

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 at least some 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 10: 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 using a training set of 18 games. All values are significant at the $p = 0.01$ level.

We expect the 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 10. 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

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.

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.

6 Performance of regime predictions in TAC SCM

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 prediction.

6.1 Prediction of regime transitions

We set the prior regime probability for the first day to 100% extreme scarcity, since all the agents start out with zero inventory on the first day. 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 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 future.

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-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 13.




Figure 13: 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. The predictions are tested across the low, medium, and high priced computer market.

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 with respect to the number of Gaussians in the GMM.

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

Table 1: Average number of regime changes and standard deviation for the testing set.

6.2 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 to determine the closeness of all individual predicted regime probabilities to the actual ones is called the *Kullback-Leibler (KL) divergence* [23, 22]. This is a quantity which measures the difference between two probability distributions in bits, meaning that 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:

$$KL(P(R)^{pred}||P(R)^{actual}) = \sum_{r \in \mathcal{R}} P(R)^{pred} \log \left(\frac{P(R)^{pred}}{P(R)^{actual}} \right) \quad (21)$$

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 [31], like: “Were yesterday’s bids all satisfied?” If the difference between the two distributions is 0 than the predictions are optimal in sense that all the probabilistic information about pricing behavior is accurate (e.g. the predicted and actual distributions match).




Figure 14: 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 are 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 a GMM with 25 components.

In Figure 14 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 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 predictions of the exponential smoother are approximately 5.6 and 45 times worse than the Markov predictions. It is typically acceptable to have 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 better for the first three days (tactical decision making) and the GMM with 16 components fits the actual regime probabilities 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 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 13 the opposite is true. Here actually the GMM with 25 components has a higher regime change prediction accuracy. The regime change success rate is based on discrete regime identification, but the KL-divergence measures the closeness of the whole

distribution. A close match of the whole distribution is important for automated applications, e.g. dynamic pricing algorithms, that utilize the continuous distribution. A human decision maker, on the other side, might get insights from prediction of discrete future regimes, since they could translate directly into some actions on the procurement or sales side.

6.3 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. 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 15 we present the algorithm used to analyze price predictions when sampling from the learned GMM Table 2 displays the results for the fitted GMMs.

```

1 Inputs:
2   pnpavg: original normalized price density
3   numBins: number of histogram bins
4   numNP: number of np in the training set
5   GMM: learned Gaussian Mixture Model
6   maxIter: number of iterations
7 Output:
8   PredErr: the overall mean prediction error
9 Process:
10  for j = 1 until maxIter
11    pnpsamp = Monte_Carlo_Sampling(GMM, numNP)
12    Error(j) =  $\frac{|pn_{p_{avg}} - pn_{p_{samp}}|}{numBins}$ 
13  end
14  meanErr =  $\overline{Error}$ 
15  PredErr = numBins ·  $\sum$  meanErr
16  return PredErr

```

Figure 15: Prediction error algorithm.

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

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 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 no classification of the market structure (for instance, monopoly vs competitive) is needed. The historical data needed to train the model are typically kept by companies as part of their business data. Many places, such as Amazon.com or eBay.com, are willing to make data on past sales available for study.

7.1 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 the competitors.

Our approach 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 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 enables 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 aggressive selling in anticipation of an upcoming oversupply situation.

7.2 Future directions

Our approach represents 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 begun to integrate it into the overall decision making process. We are currently using regime predictions to generate price trend forecasts, in order to support strategic decision making in the sales component of our agent.

We have begun design of an algorithm that will use regime predictions for tactical decision making as well. Ultimately, we plan to combine probability information supplied by the regime model with information about possible consequences of actions to optimize decision making. In particular, we plan to use the full regime-based predicted price density to optimize tactical sales pricing by maximizing expected utility [37].

In addition, we plan to apply reinforcement learning [33] to map economic regimes to internal operational regimes and operational regimes to actions, such as procurement and production scheduling. We envision an approach that would treat an operational regime as a set of rules and

parameters that are specific to a current or forecast regime, used to guide agent decision processes. For example, short-term procurement tends to be relatively expensive, but might be justified when a scarcity situation is anticipated in the near future.

8 Appendix: Summary of notation

Symbol	Definition
np	Normalized price
$\overline{\text{np}}$	Mid-range normalized price
$\widetilde{\text{np}}$	Smoothed mid-range normalized price
$\widetilde{\text{np}}^{\text{min}}$	Smoothed minimum normalized price
$\widetilde{\text{np}}^{\text{max}}$	Smoothed maximum normalized price
α	Smoothing coefficient
$p(\text{np})$	Density of normalized price
GMM	Gaussian Mixture Model
N	Number of Gaussians in the GMM
$p(\text{np} \zeta_i)$	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
$P(\zeta_i)$	Prior probability of i -th Gaussian of the GMM
$P(\zeta_i \text{np})$	Posterior probability of the i -th Gaussian of the GMM given a normalized price np
$\vec{\eta}(\text{np})$	N -dimensional vector of posterior probabilities, $P(\zeta_i \text{np})$, of the GMM components
M	Number of regimes
R_k	k -th Regime, $k = 1, \dots, M$
$\mathbf{P}(\zeta r)$	Conditional probability matrix (N rows and M columns) resulting from k -means clustering
$p(\text{np} R_k)$	Density of the normalized price np given a regime R_k
$P(R_k \text{np})$	Probability of regime R_k given a normalized price np
d	Current time
$\mathbf{T}(r_{d+1} r_d)$	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] Robert G. Brown, Richard F. Meyer, and D. A. D’Esopo. The fundamental theorem of exponential smoothing. *Operations Research*, 9(5):673–687, 1961.

- [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. Cox. Understanding Buyer and Supplier Power: A Framework for Procurement and Supply Competence. *Journal of Supply Chain Management*, 37(2):8–15, 2001.
- [9] 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.
- [10] E. F. Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417, 1970.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] T.M. Humphrey. Marshallian Cross Diagrams and Their Uses before Alfred Marshall: The Origins of Supply and Demand Geometry. *Economic Review*, pages 3–23, 1992.
- [15] I2. Next-generation planning. http://i2.com/solution_library/ng_planning.cfm, 2006.
- [16] Michael Kearns and Luis Ortiz. The Penn-Lehman Automated Trading Project. *IEEE Intelligent Systems*, pages 22–31, 2003.
- [17] Jeffrey O. Kephart, James E. Hanson, and Amy R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, 2000.
- [18] Wolfgang Ketter. *Identification and Prediction of Economic Regimes to Guide Decision Making in Multi-Agent Marketplaces*. PhD thesis, University of Minnesota, Twin-Cities, USA, January 2007.
- [19] 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.

- [20] 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.
- [21] 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 Automated Planning and Scheduling*, pages 384–392, Whistler, BC, Canada, June 2004.
- [22] Solomon Kullback. *Information Theory and Statistics*. Dover Publications, New York, 1959.
- [23] Solomon Kullback and Richard A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.
- [24] Richard Lawrence. A machine learning approach to optimal bid pricing. In *8th INFORMS Computing Society Conf. on Optimization and Computation in the Network Era*, pages 1–22, Arizona, January 2003.
- [25] Bill Mark and Raymond C. Perrault. Calo: Cognitive assistant that learns and organizes. <http://www.ai.sri.com/project/CALO>, 2006.
- [26] Cade Massey and George Wu. Detecting regime shifts: The causes of under- and overestimation. *Management Science*, 51(6):932–947, 2005.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] Koen Pauwels and Dominique Hanssens. Windows of Change in Mature Markets. In *European Marketing Academy Conf.*, Braga, Portugal, May 2002.
- [31] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423 and 623–656, July and October 1948.
- [32] Eric Sodomka, John Collins, and Maria Gini. Efficient statistical methods for evaluating trading agent performance. In *Proc. of the Twenty-Second National Conference on Artificial Intelligence*, pages 770–775, 2007.
- [33] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning - An Introduction*. The MIT Press, Cambridge, 1998.
- [34] D. Titterton, A. Smith, and U. Makov. *Statistical Analysis of Finite Mixture Distributions*. Wiley, New York, 1985.
- [35] Eric van Heck and Peter Vervest. Smart business networks: how the network wins. *Commun. ACM*, 50(6):28–37, 2007.

- [36] Peter Vervest, Eric van Heck, Kenneth Preiss, and Louis-Francois Pau. *Smart Business Networks*. Springer Verlag, Berlin, 2005.
- [37] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, Princeton, N.J., 1947.
- [38] Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, and Yevgeniy Vorobeychik. Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research*, 21:19–36, 2004.
- [39] 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.