# Identifying and Forecasting Economic Regimes in TAC SCM

**Wolfgang Ketter, John Collins, Maria Gini, Alok Gupta⋆, and Paul Schrater**
Department of Computer Science and Engineering
⋆Department of Information and Decision Sciences
University of Minnesota
{ketter,jcollins,gini,schrater}@cs.umn.edu, agupta@csom.umn.edu

## Abstract

We present methods for an autonomous agent to identify dominant market conditions, such as over-supply or scarcity, and to predict market changes. The characteristics of economic regimes are learned from historic data and used, together with real-time observable information, to identify the current market regime and to forecast market changes. We use a Gaussian Mixture Model to represent the probabilities of market prices. By clustering these probabilities we identify different economic regimes. We show that the regimes so identified have properties that represent different prevailing market conditions. We then present preliminary work on methods to predict future regime transitions. A trading agent can use these predictions to make operational and strategic decisions regarding pricing, raw material acquisition, and production. We validate our method by presenting experimental results obtained with data from the Trading Agent Competition for Supply Chain Management.

## 1 Introduction

Electronic marketplaces are gaining popularity among producers seeking to streamline their supply chains and consumers looking for opportunities. Intelligent software agents can significantly facilitate human decision processes either by helping users to select strategies or by making autonomous choices that are consistent with the human decision maker's preferences. Regardless of the application, agents need methods for making autonomous decisions under uncertain and dynamic market conditions.

The Supply-Chain Management Trading Agent Competition [Sadeh *et al.*, 2003] (TAC SCM) involves a Supply Chain Management scenario in which in each round six autonomous agents attempt to maximize profit by selling personal computers they assemble from parts they buy from suppliers. The agent with the highest bank balance at the end of the game wins. An agent has to make many decisions, such as how many parts to buy, when to get them delivered, what types of computers to build, when to sell them, and at what price. Availability of parts and demand for computers varies randomly through the game and across three market segments (low, medium, and high computer price). The market segments are affected not only by random variations in supply and demand, but also by the actions of the other agents. The small number of agents and their ability to adapt, to change strategy during the game, and to manipulate the market makes the game highly dynamic and uncertain.

The problem we address in this paper is how market conditions, such as oversupply or scarcity, can be detected and exploited by an autonomous agent. The long term objective of our work is to show how knowledge of current and anticipated market conditions can enable an agent to make better operational and strategic decisions, such as how many parts to order, how to schedule production, and how to price its products. While this type of prediction about the economic environment is commonly used at the macro economic level, such predictions are rarely done for a micro economic environment.

We first introduce methods to identify distinguishable economic conditions, which we call *regimes*, such as over-supply or scarcity of products. This is accomplished using historical data from previous games and real-time data available during the game. We then present preliminary ideas on how to forecast regime transitions. A short review of related work concludes the paper.

## 2 Economic Regime Identification

We believe that market conditions can be characterized by statistical patterns, and that such patterns can be learned offline from historical data. We call those distinguishable market conditions *regimes*. Since supply and demand in TAC SCM change in each of the market segments independently of the other segments, we will limit our considerations to a market segment, and assume the same methods are applicable to each market segment.

To give an intuition of how prices for the same type of computer change during the game, we show in Figure 1 the probability of receiving an order for a given offer price for computers of type 1 during one of the games played in the finals of TAC SCM 2004. We can see how the slope of the curve and its position change 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.

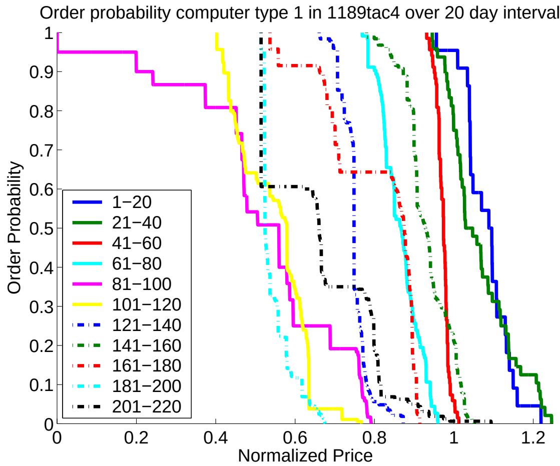


Figure 1: Game 1189@tac4 (Final TAC SCM 04) – Probability of order for computer type 1 by offer price. The plot shows the curves every 20 days during the game.

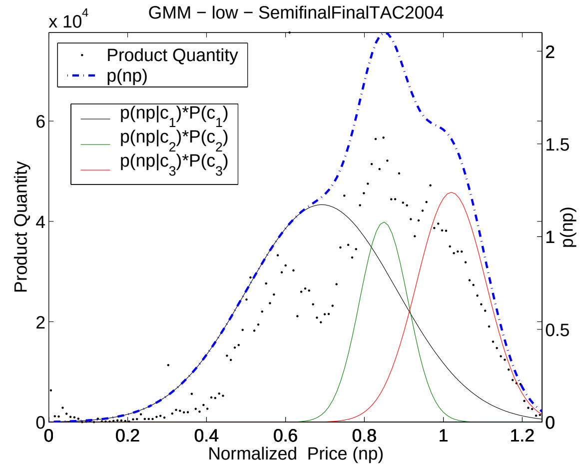


Figure 2: The Gaussian mixture model for the low market segment. Data are from 26 games from finals and semifinals of TAC SCM 2004.

Clearly the market conditions change during the game, and this should affect the strategy of the agent. When there is scarcity, prices are higher, so the agent can price aggressively. In balanced situations prices are lower and have more spread, so the agent has a range of options for maximizing expected profit. In over-supply situations prices are lower. The agent should primarily control costs, and therefore either do pricing based on costs, or wait for better market conditions.

## 2.1 Off-line analysis of data

In our approach, we identify and characterize market regimes by analyzing off-line data from previous games. The agent can then use these results along with real-time observable information to identify regimes during the game, forecast regime transitions, and adapt its procurement, production, and pricing strategy accordingly.

For our experiments, we used data from a set of 26 games played during the semi-finals and finals of TAC SCM 2004. The number of games played was 30, but we left out the games where some computers were sold for \$0. The mix of players changed from game to game, the total number of players was 12 in the semifinals and 6 in the finals.

We define regimes with the help of a Gaussian mixture model (GMM).

Each computer type has a different nominal price, which is the sum of the nominal cost of each of the parts needed to build it. We normalize the prices across the different computer types in each market segment. We call $\mathrm{np}$ the normalized price. We apply the EM-Algorithm [Dempster *et al.*, 1977] to determine the Gaussian components of the GMM, $N[\mu_i, \sigma_i](\mathrm{np})$, and their prior probability, $P(c_i)$. The density of the normalized price can be written as:

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

where $p(\mathrm{np}|c_i)$ is the $i$-th Gaussian from the GMM. An example of the Gaussians is shown in Figure 2. For our experiments we chose $N = 3$, because we found experimentally that this provides a good balance between quality of approximation and simplicity of processing.

Using Bayes' rule we determine the posterior probability:

$$P(c_i|\mathrm{np}) = \frac{p(\mathrm{np}|c_i)\, P(c_i)}{\sum_{i=1}^{N} p(\mathrm{np}|c_i)\, P(c_i)} \qquad \forall i = 1, \cdots, N \quad (2)$$

We then define the N-dimensional vector, whose components are the posterior probabilities from the GMM,

$$\vec{\eta}(\mathrm{np}) = [P(c_1|\mathrm{np}), P(c_2|\mathrm{np}), \ldots, P(c_N|\mathrm{np})] \qquad (3)$$

and for each normalized price $\mathrm{np}_j$ we compute $\vec{\eta}(\mathrm{np}_j)$ which is $\vec{\eta}$ evaluated at the $\mathrm{np}_j$ price. We cluster these collections of vectors using k-means. The center of each cluster corresponds to regime $R_k$ for $k = 1, \cdots, M$, where $M$ is the number of regimes.

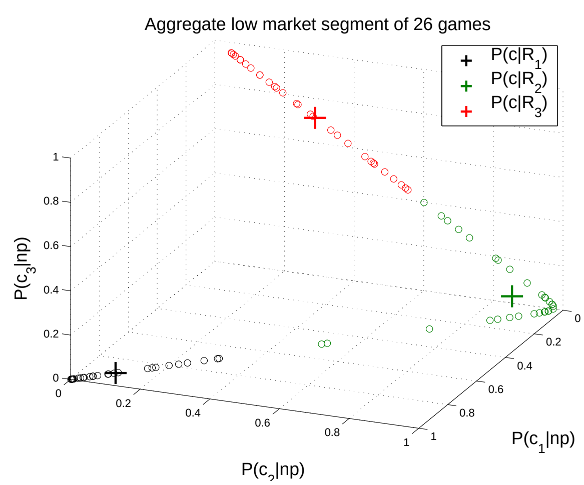


Figure 3: K-means clustering applied to the posterior probability $P(c|\mathrm{np})$ in the low market segment.

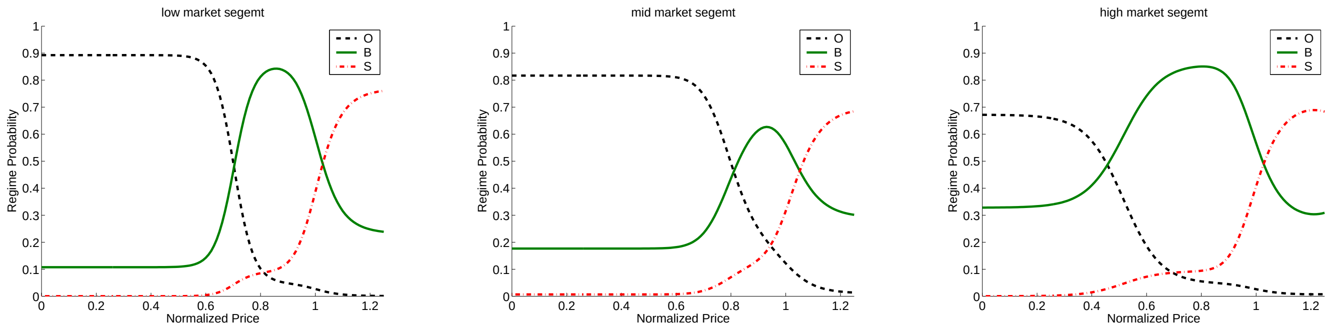


Figure 4: Regime probabilities over normalized price for the low (left), medium (middle) and high (right) market segments. These are computed off-line from 26 games.

Figure 3 shows the cluster centers, which correspond to regimes, for the low market segment. The figure shows only some of sample points for better visualization.

We distinguish three regimes, namely over-supply ($R_1$), balanced ($R_2$), and scarcity ($R_3$). Regime $R_1$ represents a situation where there is a glut in the market, i.e. an over-supply situation, which depresses prices. Regime $R_2$ represents a balanced market situation, where most of the demand is satisfied. In regime $R_2$ the agent has a range of options of price vs sales volume. Regime $R_3$ represents a situation where there is scarcity of products in the market, which increases prices. In this case the agent should price close to the customer reserve price – the maximum price a customer is willing to pay.

The number of regimes was selected a priori, after examining the data and looking at economic analyses of market situations. Both the computation of the GMM and k-means clustering were tried with different initial conditions, but consistently converged to the same results. Correlation analysis (see, later, Figure 8) shows that regimes can be characterized in terms of market quantities, such as prices and ratio of offer to demand.

We can rewrite $p(\mathrm{np}|c_i)$ in a form that shows the dependence of the normalized price $\mathrm{np}$ not on the Gaussian $c_i$ of the GMM, but on the regime $R_k$:

$$P(\mathrm{np}|R_k) = \sum_{i=1}^{N} p(\mathrm{np}|c_i)\, P(c_i|R_k). \quad (4)$$

The probability of regime $R_k$ dependent on the normalized price $\mathrm{np}$ can be computed using Bayes rule as:

$$P(R_k|\mathrm{np}) = \frac{P(\mathrm{np}|R_k)\, P(R_k)}{\sum_{k=1}^{M} P(\mathrm{np}|R_k)\, P(R_k)} \quad \forall k = 1, \cdots, M. \quad (5)$$

where M is the number of regimes, which in our case is 3. The prior probabilities $P(R_k)$ of the different regimes are determined by a counting process over multiple games.

Figure 4 depicts the regime probabilities for the three market segments. Each regime is clearly dominant over a range of normalized prices. To make things more intuitive, we label regime $R_1$ as $O$ for over-supply, regime $R_2$ as $B$ for balanced, and regime $R_3$ as $S$ for scarcity. The relative dominance and range of the different regimes varies among the market segments, but we can see, as expected, that oversupply corresponds to lower prices, a balanced situation to prices closer to the average, and scarcity to high prices. We assume this reflects different agent pricing and inventory-management strategies. The high market segment offers higher profit per computer, hence the balanced regime extends over a larger range of normalized prices. In the low market segment the profit per computer is low, hence the balanced regime extends over a much smaller range of normalized prices.

## 2.2 Online identification of current regime

During the game, the agent can estimate every day the current regime by calculating the mean normalized price $\overline{\mathrm{np}}_{day}$ for the day and by selecting the regime which has the highest probability, i.e. $\mathrm{argmax}_{1 \le k \le M}\, \vec{P}(R_k|\overline{\mathrm{np}}_{day})$.

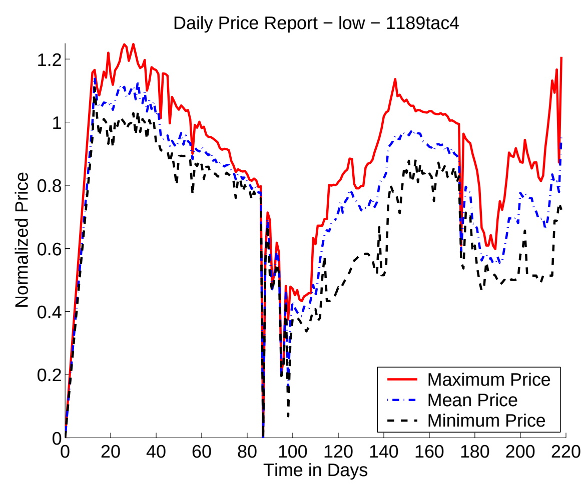


Figure 5: Game 1189@tac4 (Final TAC SCM 04) – Minimum and maximum daily prices of computers sold, as reported during the game, and mean price. The mean price is computed after the game using the game data.

Unfortunately, the agent has only limited market information. The best it can do is to estimate the mean of the normalized price of the computers sold the previous day using the daily price reports of minimum and maximum prices. This is only an estimate, since the quantity if computers sold is not known. An example that shows how the mean differs from the estimate is shown in Figure 5.

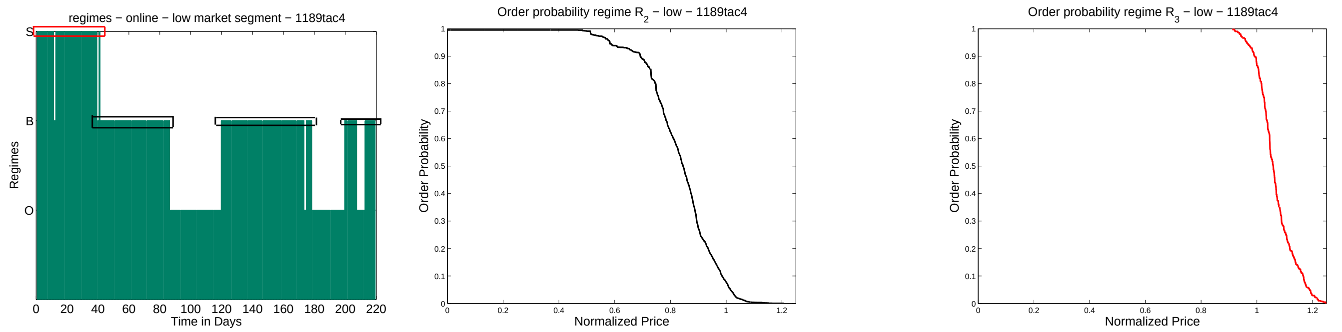


Figure 6: Game 1189@tac4 (Final TAC SCM 04) – Regimes over time for the low market computed online every day (left), probability of receiving an order by normalized price for a balanced situation ($R_2$ indicated by $B$) (middle) and for a scarcity situation ($R_3$ indicated by $S$) (right).

This estimate can then be used to identify the corresponding regime online, as shown in Figure 6 (left). The data are from game 1189@tac4, which was not in the training set of games used to develop the regime definitions. The middle and right parts of Figure 6 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.

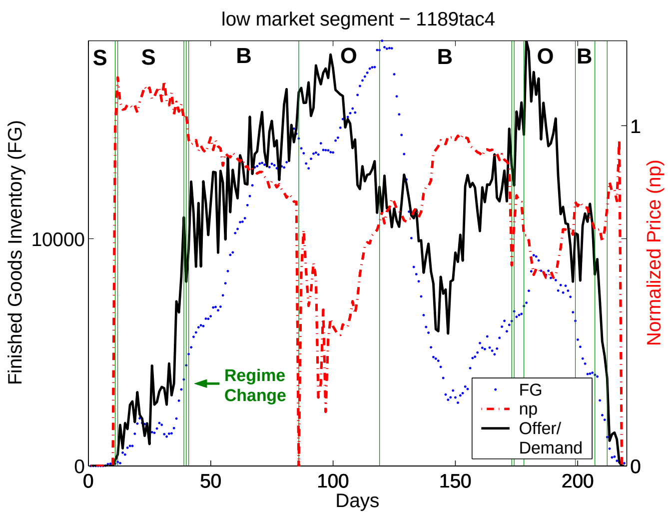


Figure 7: Game 1189@tac4 (Final TAC SCM 04) – Relationships in the low market between regimes and normalized prices. We show the ratio of offer to demand, and the available finished goods inventory. The dominant regimes are labeled along the top. On day 141 the ratio of offer to demand is about 1.4 and on day 179 it is about 4.5.

Figure 7 shows the quantity of the finished goods inventory ($FG$), the ratio of offer to demand, and the normalized price over time. The ratio of offer to demand represents the proportion of the market that is satisfied. These factors clearly correlate with market regimes, but they are not directly visible to the agent during the game. The regimes identified by our approach are superimposed, where $S$ (or $R_3$) represents scarcity, $B$ (or $R_2$) balanced, and $O$ (or $R_1$) over-supply.

Economic intuition tells us that when there is over-supply prices are low and vice versa. Figure 7 shows a clear correlation between regimes and market parameters, substantiating our economic regime characterization. For example, the figure shows that when the offer to demand ratio is high (i.e. over-supply) prices are low and vice versa.

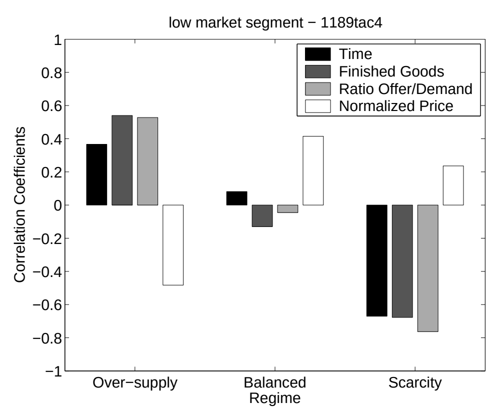


Figure 8: Game 1189@tac4 (Final TAC SCM 04) – Correlation coefficients between regimes and time in the game, the ratio of offer to demand, normalized price (np), and quantity of finished goods inventory in the low market segment.

A correlation analysis of the market parameters is shown in Figure 8. The p-values for the correlation analysis are all less than 0.01. Regime $R_1$ (over-supply) correlates strongly and positively with time, ratio of offer to demand, and quantity of finished goods inventory, and negatively with normalized price. On the other hand, in Regime $R_3$ (scarcity) we observe a strong negative correlation with the selected market parameters.

Figure 9 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

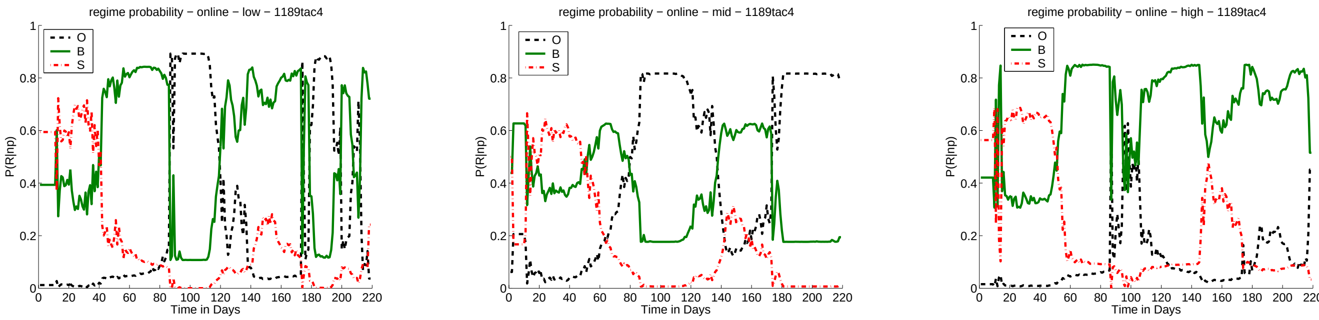


Figure 9: Game 1189@tac4 (Final TAC SCM 04) – Regime probabilities over time computed online every day for the low (left), medium (middle) and high (right) market segment.

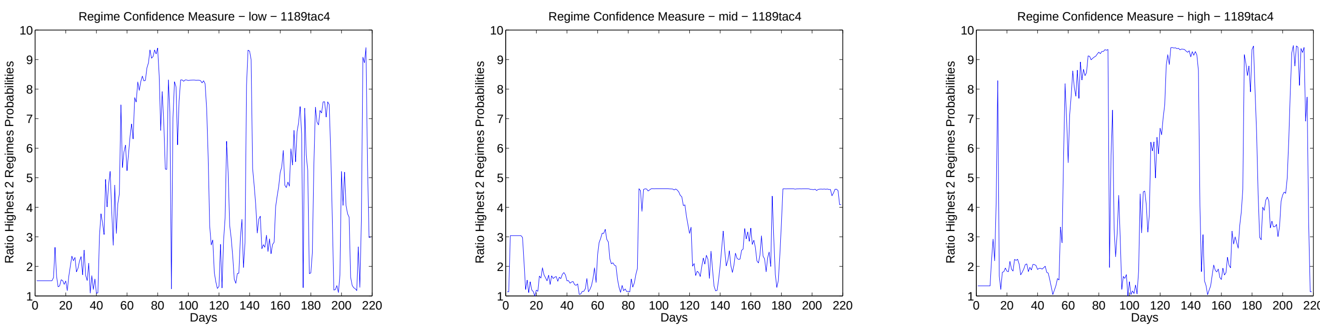


Figure 10: Game 1189@tac4 (Final TAC SCM 04) – Ratio of probabilities of best and second best regime computed every day for the low (left), medium (middle), and high (right) market segment. A high ratio corresponds to a high confidence in the current regime and a ratio close to one indicates a mixture of almost equally likely regimes.

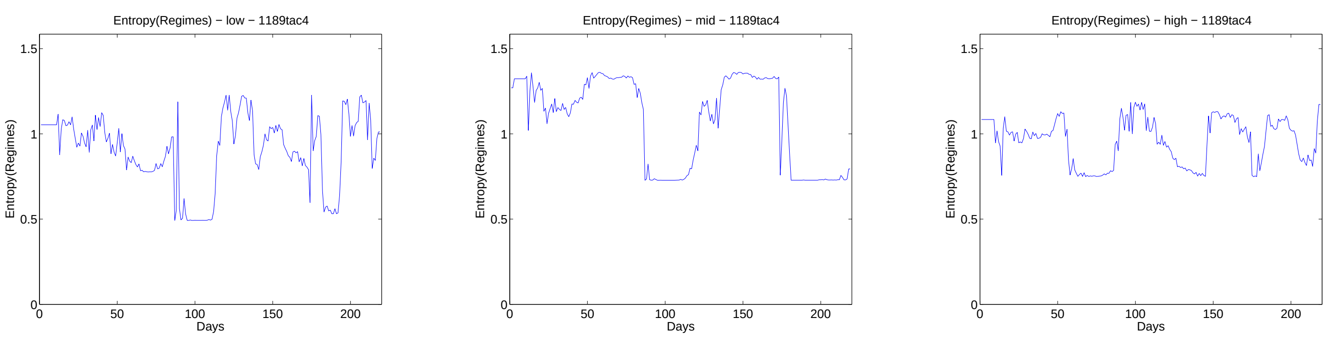


Figure 11: Game 1189@tac4 (Final TAC SCM 04) – Daily entropy values of the 3 regimes for the low (left), medium (middle), and high (right) market segment.

which strategy, or mixture of strategies, to follow.

To estimate the confidence in the regime identification an agent can use the ratio of the probability of the highest vs the second highest regime. A high ratio corresponds to high confidence in regime identification, while a ratio close to one indicates that the current market situation is a mixture of almost equally likely regimes. This can be an indication of an upcoming regime change. Examples for the three market segments in game 1189tac4 are shown in Figure 10.

We can observe that the ratios are smaller in the medium market segment than in the other two markets. We believe that this can be in part attributed to the fact that computers in the medium segment are built out of parts that can also be used to build computers in either the low or the high market segment. Agents might decide to switch some of their production to the market segment that appears most profitable, so making the regimes in the medium market not as dominant.

An additional measure of the confidence in the regime identification is the entropy of the set $S$ of probabilities of the regimes given the mean of the daily normalized price $\overline{\mathrm{np}}_{day}$, where

$$S = \{P(R_1|\overline{\mathrm{np}}_{day}), \cdots, P(R_M|\overline{\mathrm{np}}_{day})\}$$

and

$$\text{Entropy}(S) \equiv \sum_{k=1}^{M} -P(R_k|\overline{\text{np}}_{day}) \log_2 P(R_k|\overline{\text{np}}_{day}). \quad (6)$$

This confidence measure based on entropy has the advantage that it includes all regimes as opposed to the ratio confidence measure which only includes the regimes with the highest and the second highest probability. 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. Examples for the three market segments in game 1189@tac4 are shown in Figure 11.

## 3 Regime Prediction

From the perspective of making decisions the behavior of an agent should depend on the current market regime as well as expectation of future regimes. The fixed start and end boundaries of the game have a strong impact on market regimes. Prices are higher at the beginning of the game, when computers are very scarce, and often fall toward the end of the game when every agent is trying to sell off its inventory. We believe that an agent which is capable of predicting regime changes has an advantage over its competitors.

We model the regime estimation process as a discrete, semi-Markov process [Levinson, 1986]. In a semi-Markov process the future depends not only on the present state, but also on the length of time the process has spent in the state. Our preliminary observations, led us to believe that the market exhibits time dependent behavior. For instance, in Figure 6 (left) we see long stable regimes and rapid switches to new regimes.

In order to model this behavior we construct a Markov transition matrix, $\mathbf{T}_{\text{predict}}$, which is a weighted sum of two matrices, the steady state matrix $\mathbf{T}_{\text{steady}}$ and the change matrix $\mathbf{T}_{\text{change}}$. $\mathbf{T}_{\text{steady}}$ is the $M \times M$ identity matrix, where $M$ is the number of regimes. $\mathbf{T}_{\text{change}}$ represents the posterior probability of transitioning to a regime given the current regime. $\mathbf{T}_{\text{change}}$ is computed off-line by a counting process.

$$\begin{aligned} \mathbf{T}_{\text{predict}}(r_{t+1}|r_t) &= (1-\omega(.)) \times \mathbf{T}_{\text{steady}}(r_{t+1}|r_t) \\ &\quad + \omega(.) \times \mathbf{T}_{\text{change}}(r_{t+1}|r_t) \end{aligned} \quad (7)$$

where $\omega(.)$ represents the probability of a regime change, and $r_t$ represents the current regime.

To compute the value of $\omega(.)$, we need to introduce a few variables. We model the time $\tau_i$ spent in regime $R_i$ before the transition to regime $R_j$ occurs as a random variable with distribution $F_{ij}$. $\tau_i$ is measured off-line.

To model $F_{ij}$ we hypothesized that the probability density of $\tau_i$ is dependent on the current regime, $R_i$, i.e. $p(\tau_i|R_i)$. We computed the frequency of all values of $\tau_i$ in ascending order, shown in Figure 12, and fitted different distributions. The Gamma distribution, $g(t;\alpha,\lambda)$ (Equation 8) is a reasonable fit to the data, as shown in Figure 13.

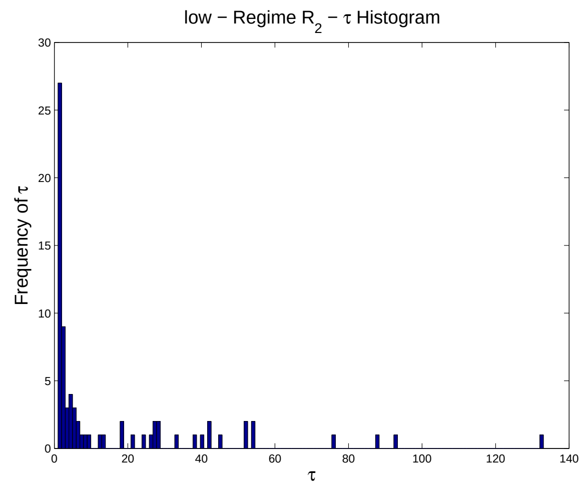


Figure 12: $\tau$ values for regime $R_2$ (balanced) for the low market segment shown in ascending order. The $\tau$ values were collected from 10 games from the TAC SCM 04 semifinal and final games.

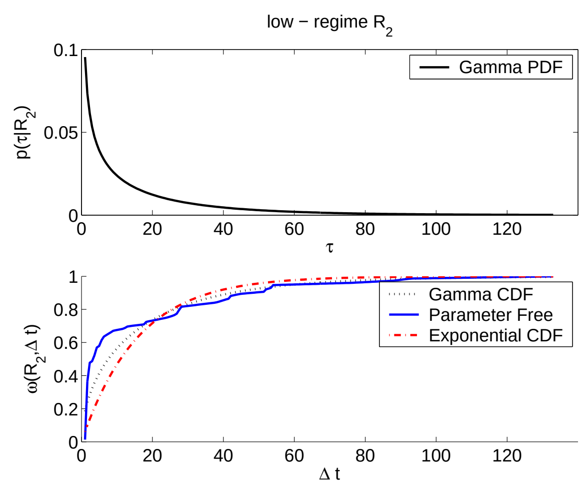


Figure 13: Fitted Gamma pdf for $p(\tau|R_2)$ (top); Cumulative distributions: $\omega(r = R_2, \Delta t)$ is the probability of transitioning out of regime $R_2$, $\Delta t$ is the elapsed time since the last regime change (bottom). Data are for the low market segment.

The gamma density function, $g(t;\alpha,\lambda)$, depends on two parameters, $\alpha$ and $\lambda$:

$$g(t;\alpha,\lambda) = \begin{cases} \frac{\lambda^\alpha}{\Gamma(\alpha)} t^{\alpha-1} e^{-\lambda t} & \text{for } t \geq 0 \\ 0 & \text{for } t < 0 \end{cases}, \quad (8)$$

where $\Gamma(x)$ is the gamma function, which is defined as

$$\Gamma(x) = \int_0^\infty u^{x-1} e^{-u} \, \text{d}u \quad x \geq 0, \alpha > 0, \lambda > 0.$$

The parameters were fitted separately for each regime using a maximum likelihood procedure. After applying the fitting

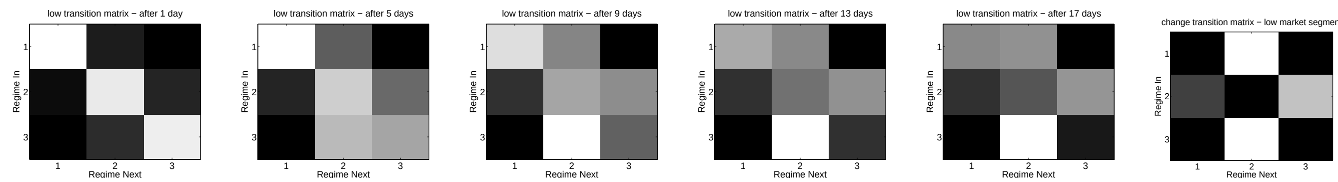


Figure 14: $\mathbf{T}_{\text{predict}}$ Markov transition matrices as a function of time for the low market segment. $\mathbf{T}_{\text{predict}}$ is a time-varying combination of $\mathbf{T}_{\text{steady}}$ and $\mathbf{T}_{\text{change}}$ for a regime with an average value of $\tau = 15$. The $\tau$ values are based on 10 games from the TAC SCM semifinal and final 2004. The transition matrix probabilities are shown as shades of gray, with white representing 1 and black representing 0. Rows are for the current regime and columns for the next regime. $\mathbf{T}_{\text{predict}}$ is evaluated for 1, 5, 9, 13 and 17 days from left to right. The rightmost matrix represents $\mathbf{T}_{\text{change}}$ without mixing.

procedure we obtained $\alpha = 0.5287$ and $\lambda = 0.0333$ for regime $R_2$ in the low market segment.

We define $\Delta t$ as the time since the last regime transition at $t_0$:

$$\Delta t = t - t_0 \tag{9}$$

The probability of a regime transition $\omega(r, \Delta t)$, shown in Figure 13 (bottom), from the current regime, $r$, with respect to the time $\Delta t$ that has elapsed since the last regime transition, $t_0$, is given by:

$$\omega(r = R_i, \Delta t) = \int_0^{\Delta t} p(\Delta t | r = R_i)\, \mathrm{d}\Delta t \tag{10}$$

where $p(\Delta t | r = R_i) = g(\Delta t; \alpha_i, \lambda_i)$, and the parameters $\alpha_i, \lambda_i$ are fitted separately for each regime.

Figure 14 shows the time dependent Markov prediction matrices for the low market segment for 1, 5, 9, 13, and 17 days after a regime shift occurred. The rightmost figure is the change matrix for the same market segment. We observe that as time passes $\mathbf{T}_{\text{predict}}$ becomes more similar to $\mathbf{T}_{\text{change}}$.

## 4 Related Work

Marketing research methods have been developed to understand the conditions for growth in performance and the role that marketing actions play to improve sales. For instance, in [Pauwels and Hanssens, 2002], an analysis is presented on how in mature economic markets strategic windows of change alternate with long periods of stability.

Model selection is the task of choosing a model of optimal complexity for the given data. A good overview of concepts, theory and model selection methods is given in [Cherkassky and Mulier, 1998].

Much work has focused on models for rational decision-making in autonomous agents. Ng and Russel [2000] 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 TAC SCM.

Carmel and Markovitch [1993] describe a game-player that tries to analyze and learn the strategy of its opponent. They discuss the benefits of using a model of the opponent strategy, and give an algorithm called M* (a generalization of the standard minimax algorithm) that attempts to exploit the opponent strategy. M* assumes that the opponent's search depth and evaluation function are known.

Chajewska, Koller, and Ormoneit [2001] show 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 the maximization of the expected utility [von Neumann and Morgenstern, 1947]. 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 price records [Pardoe and Stone, 2004], or by estimating the relationship between offer price and order probability with a linear cumulative density function (CDF) [Benisch *et al.*, 2004], or by using a reverse CDF and factors such as quantity and due date [Ketter *et al.*, 2004], or by letting other agents set the price and trying to follow [Dahlgren and Wurman, 2004].

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 which defines the relationship between dependent and independent variables has the same structure. However, as shown in Figure 1, these functional relationships have a different structure for different regimes. Therefore, a clustering based approach that does not assume a functional relationship maybe the best way to identify a regime.

Wellman et al. [2005] 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 to guide the operational and strategic behavior of the agent.

## 5 Conclusions and Future Work

We have presented an approach to characterizing and predicting economic market conditions. Our approach recognizes that different market situations have qualitative differences that can be used to guide the strategic and tactical behavior of an agent. Unlike regression-based methods that try to predict prices directly from demand and other observable factors, our approach recognizes that prices are also influenced by non-observable factors, such as the inventory positions of the other agents. Unlike price-following methods, our approach promises to enable an agent to anticipate and prepare for regime changes, for example by building up inventory in anticipation of better prices in the future or by selling in anticipation of an upcoming oversupply situation.

We have demonstrated the effectiveness of our approach by characterizing the market conditions in games played in the semifinals and finals from TAC SCM 2004.

Our next step is to complete the prediction of future regimes, to design and evaluate sales strategies that take advantage of regime prediction, and to integrate them into the decision making process of our agent. This, for instance, will prevent the agent from purchasing parts when there is an over-supply of computers in the market, and encourage buying parts when under-supply is anticipated. We believe that our proposed formulation will allow the agent to operate effectively on a daily basis as well as to engage in strategic pricing and market manipulation.

### Acknowledgements

Partial funding is acknowledged from NSF under grant EIA-0324864.

## References

[Benisch *et al.*, 2004] 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 Computational Logic*, 4(3):29–37, 2004.

[Carmel and Markovitch, 1993] D. Carmel and S. Markovitch. Learning models of opponent's strategy in game playing. Technical report, Technion-Israel Institute of Technology, 1993.

[Chajewska *et al.*, 2001] 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.

[Cherkassky and Mulier, 1998] Vladimir Cherkassky and Filip Mulier. *Learning from data - Concepts, Theory, and Methods*. John Wiley & Sons, INC., New York, 1998.

[Dahlgren and Wurman, 2004] Erik Dahlgren and Peter Wurman. PackaTAC: A conservative trading agent. *SIGecom Exchanges*, 4(3):33–40, 2004.

[Dempster *et al.*, 1977] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, Series B, 39(1):1–38, 1977.

[Ketter *et al.*, 2004] Wolfgang Ketter, Elena Kryzhnyaya, Steven Damer, Colin McMillen, Amrudin Agovic, John Collins, and Maria Gini. MinneTAC sales strategies for supply chain TAC. In *Proc. of the Third Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1372–1373, New York, July 2004.

[Levinson, 1986] S E Levinson. Continuously variable duration hidden markov models for automatic speech recognition. *Comput. Speech Lang.*, 1(1):29–45, 1986.

[Ng and Russell, 2000] 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.

[Pardoe and Stone, 2004] 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.

[Pauwels and Hanssens, 2002] Koen Pauwels and Dominique Hanssens. Windows of Change in Mature Markets. In *European Marketing Academy Conf.*, Braga, Protugal, May 2002.

[Sadeh *et al.*, 2003] Norman Sadeh, Raghu Arunachalam, Joakim Eriksson, Niclas Finne, and Sverker Janson. TAC-03: A supply-chain trading competition. *AI Magazine*, 24(1):92–94, 2003.

[von Neumann and Morgenstern, 1947] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, Princeton, N.J., 1947.

[Wellman *et al.*, 2005] Michael P. Wellman, Joshua Estelle, Satinder Singh, Yevgeniy Vorobeychik, Christopher Kiekintveld, and Vishal Soni. Strategic interactions in a supply chain game. *Computational Intelligence*, 21(1):1–26, 2005.