

Economic Regime Identification and Prediction in TAC SCM Using Sales and Procurement Information

Frederik Hogenboom

Erasmus Sch. of Econ.
Erasmus University Rotterdam
fhogenboom@ese.eur.nl

Uzay Kaymak

Erasmus Sch. of Econ.
Erasmus University Rotterdam
kaymak@ese.eur.nl

Wolfgang Ketter

Rotterdam Sch. of Mgmt.
Erasmus University Rotterdam
wketter@rsm.nl

John Collins

Computer Science and Engr.
University of Minnesota
jcollins@cs.umn.edu

Jan van Dalen

Rotterdam Sch. of Mgmt.
Erasmus University Rotterdam
jdalen@rsm.nl

Alok Gupta

Carlson Sch. of Mgmt.
University of Minnesota
agupta@csom.umn.edu

Abstract

Our research is focused on the effects of the addition of procurement information (offer prices) to a sales-based economic regime model. This model is used for strategic, tactical, and operational decision making in dynamic supply chains. We evaluate the performance of the regime model through experiments with the MinneTAC trading agent, which competes in the TAC SCM game. The new regime model has an overall predictive performance which is equal to the performance of the existing model. Regime switches are predicted more accurately, whereas the prediction accuracy of dominant regimes does not improve. However, because procurement information has been added to the model, the model has been enriched, which gives new opportunities for applications in the procurement market, such as procurement reserve pricing.

Introduction

Nowadays, markets are extremely competitive, and thus it is important to gain insight into the dynamics of supply chains and to research supply chain optimization possibilities, both for individual elements in the chain, as well as for the chain as a whole. For instance, in correctly predicting future market conditions, i.e., economic regimes (Ketter 2007), lie competitive advantages, e.g., one can anticipate on upcoming scarcities in the sales market by adjusting procurement policies and sales prices in advance. This can save for instance storage costs, and increase profits. Thus, one could benefit greatly from being able to make tactical and strategic decisions in an uncertain market, based on identified and predicted economic regimes. Combining techniques from computer science with economic theory to solve problems in economic environments contributes to novel approaches to existing problems.

Ketter introduced an economic regime model, which is based on sales information (Ketter 2007). This model can be applied to any real or simulated market situation. However, procurement information has not been

used so far for identifying and predicting regimes. This information is truly valuable for determining economic regimes, since it captures specific market characteristics. For instance, an increase in the amount of components sold in the procurement market could indicate an expected scarcity in the sales market, as manufacturers are building stocks.

We present an extension to the regime model as introduced by Ketter, which is implemented in the MinneTAC trading agent (Collins, Ketter, and Gini 2009). We investigate the effects of adding procurement information to this model. The performance of the regime model is evaluated through experiments on the quality of regime probability predictions, and checking correlations with existing market conditions.

The MinneTAC trading agent has competed for several years in the Trading Agent Competition for Supply Chain Management (TAC SCM). TAC SCM is an annual international competition for designing trading agents for a simulated personal computer supply chain (Collins et al. 2005). The supply chain includes customers, traders, and suppliers and the computer market is divided into market segments. Traders have to procure components from suppliers and sell assembled computers to customers. Only limited information is available to traders, such as yesterday's maximum and minimum market prices, which makes predicting the market a non-trivial task.

TAC SCM has attracted researchers from all over the world, because its environment is designed in such a way that it contains many characteristics that can be found in real-life supply chains, such as the behavior of unpredictable opponents and interdependent chain entities. The simulated supply chain of the TAC SCM game offers many research opportunities into various subjects, such as price setting strategies and prediction strategies for competitor behavior or market characteristics and developments.

The paper is organized as follows. First, we continue with introducing the existing economic sales-based regime model. Subsequently, we define and evaluate a new economic regime framework based on both

sales and procurement information. Finally, conclusions are drawn and future work is suggested.

Background and Related Work

Regime identification and prediction using the economic regime model are used for making sales decisions in the MinneTAC agent (Ketter et al. 2009). Every day, for each individual product, the probabilities of current and future economic regimes are determined, which are used for forecasting price densities. With the help of the current and future price density we are able to predict market prices, market price trends, and the customer acceptance probability for specific offers. Regimes are used for both tactical and strategic decisions, such as product pricing and production planning.

The regimes in the TAC SCM game can be considered as a set of characteristics which apply to a certain period of days. The identification and prediction of regimes is done so that different behavior can be modeled for different situations, which is also referred to as a switching model. The agent's problems can be solved differently depending on the state of its environment, which influences the accuracy of the agent's predictions and the amount of profit made.

Regimes are used in multiple contexts, such as political and economic contexts. In general, a regime refers to a set of conditions. In economic context, regimes are also referred to as business cycle phases. These phases are commonly used in macro-economic environments, as is the case in (Osborn and Sensier 2002). However, in (Ketter 2007), regimes are applied in the micro-economic environment of the TAC SCM game. This makes sense, since an economic environment is simulated and one can capture (economical) characteristics in economic regimes, enabling an agent to reason based on certain market conditions.

Other applications of regimes can be found in electricity markets. These markets often are oligopolies, which is also the case in TAC SCM. Becker et al. and Mount et al. both model spikes in electricity prices as a regimes switching model, based on Markov switching models (Becker, Hurn, and Pavlov 2007; Mount, Ning, and Cai 2006).

Over the past decades, research has been into identifying and predicting regimes, but also into regime changes. Regime changes are important events in time series, in which one can obtain strategic advantage if they are identified or predicted correctly. Recently, Massey and Wu (2005) emphasize this importance of the ability to detect and respond to regime shifts, as it is critical to economic success. If these shifts are not detected, this could lead to lower profits. Also, Massey et al. elaborate on the causes of under- and overreaction to (predicted) regime shifts.

The foundation of research into regime shifts lies in 1989, when Hamilton published a paper about modeling regime changes using postwar U.S. real GNP as input (Hamilton 1989). Hamilton used Markov matrices

to observe regime shifts, by drawing probabilistic inference about whether and when they may have occurred based on the observed behavior of series.

The algorithms of the regime model implemented by the MinneTAC agent, which identifies and predicts economic regimes, are based on economic theory and incorporate some adapted techniques. The model's regimes are identified as extreme scarcity, scarcity, a balanced situation, oversupply, and extreme oversupply (Ketter 2007). We define the regime set as $R = \{ES, S, B, O, EO\}$. Each day of the game, the regime probability distribution is determined and a regime prediction is made. Thus, each regime R_k in set $R \forall k = 1, \dots, M$ (where $M = 5$) has a certain likelihood to hold as dominant regime. The regime probabilities in set R sum up to 1 and the regime with the highest probability is considered to be dominant.

Currently, the regime model identifies and predicts economic regimes based on yesterday's normalized mean (mid-range) sales price (Ketter et al. 2007) for an arbitrary day d , which is also referred to as np_{d-1} , as well as on quantities. Also, for some predictions, the entire history of normalized mean sales prices is used. The problem is that in supply chains, data for the day itself is not available, and thus identification is done based on historical information, whereas predictions are made for day d up to planning horizon h (usually twenty days).

The data the model is based on is normalized, as normalization allows for machine learning across markets, i.e., over different products, which can be compared qualitatively, and it enables whole market forecasts. Also, the range of the variable is fixed and thus known beforehand, which simplifies computations.

Regime identification is currently done by offline and online machine learning. Offer acceptance probabilities associated with given product prices (approximated using a Gaussian Mixture Model (Titterton, Smith, and Makov 1985)), derived from observable historical and current sales market data, are clustered offline using the K-Means algorithm (MacQueen 1967), which yields distinguishable statistical patterns (clusters). These clusters are labeled with the proper regimes after correlation analyses. Regime probabilities, which are indicative of how market conditions are, are determined by calculating the normalized price density of all clusters, given sales prices.

There are several techniques for predicting regimes, each of which is most suitable for a specific time span. Today's regimes can be predicted based on exponentially smoothed price predictions, as extensively elaborated in (Ketter et al. 2008). Short-term regime prediction for tactical decision making is done by using a Markov prediction process. This process is based on the last normalized smoothed mid-range price. To this end, Markov transition matrices, which are created offline by a counting process over past games, are being used. Long-term regime prediction is done by using a Markov correction-prediction process. This process is almost equal to the short-term regime prediction, but

is based on all normalized smoothed mid-range prices up and until the previous day, instead of just the last normalized smoothed mid-range price.

Extending the Regime Model

We extend the regime model as introduced by Ketter et al. (2009) with procurement information. Our model differentiates from the model introduced in (Ketter et al. 2009) in the fact that it is based on procurement information, and not solely on sales information. In order to be able to select the most promising procurement variable, we apply the information gain metric (Andrews et al. 2007) to a data set containing procurement information on prices, quantities, offers, orders, and requests for quotation (RFQs) gathered from historical game data¹. We now continue with discussing some details of the information gain metric.

Information Gain

The information gain is an entropy-based metric that indicates how much better we can predict a specific target by knowing certain features. In our case, the target is the dominant regime and the features are procurement variables. According to Mitchell (Mitchell 1997), the entropy is a commonly used measure in information theory, which characterizes the purity of an arbitrary collection of examples. Let W be a collection of game results, numW the number of possible values of W (i.e., regimes), and $P(w)$ the probability that W takes on value w . Assuming a uniform probability distribution, the latter probability is equal to the proportion of W belonging to class w . The entropy of a collection of game results, entropy(W), is defined as

$$\text{entropy}(W) = \sum_{w=1}^{\text{numW}} -P(w) \log_2 P(w) . \quad (1)$$

Now let V be an attribute (procurement variable), numV the number of possible values of V , $P(v)$ the probability that V takes on value v , and $P(w|v)$ the probability that W takes on value w , given v . The entropy of a collection of game results W given an attribute V , entropy($W|V$), is defined as

$$\text{entropy}(W|V) = \sum_{v=1}^{\text{numV}} P(v) \left(\sum_{w=1}^{\text{numW}} -P(w|v) \log_2 P(w|v) \right) . \quad (2)$$

Using (1) and (2), the information gained on outcome W from attribute V , gain(W, V), can be calculated using

$$\text{gain}(W, V) = \text{entropy}(W) - \text{entropy}(W|V) . \quad (3)$$

¹Data set contains 2007 Semi-Finals games played on the SICS tac5 server (IDs: 9321–9328), 2007 Finals games played on the SICS tac3 server (IDs: 7306–7313), 2008 Semi-Finals games played on the University of Minnesota’s (UMN) tac02 server (IDs: 761–769), and 2008 Finals games played on the UMN tac01 server (IDs: 792–800).

Here, the entropy of a collection of game results W given an attribute V is subtracted from the entropy of W . For more information on the calculation of the information gain, see (Andrews et al. 2007). Applying the information gain metric to several possible procurement variables using our data set results in Table 1. The higher the score, the more information the variable adds to the model.

Variable	Gain
Offer price	0.7393
Order price	0.5400
RFQ lead time	0.5106
RFQ reserve price	0.4909
Ratio orders / offers	0.4555
Order quantity	0.4310
RFQ quantity	0.3901
Demand	0.3833
Offer quantity	0.3174

Table 1: Information gain scores for several procurement variables.

The latter table shows that component offer prices (recalculated on a per-product basis) are most likely to improve the predictive capabilities of the regime model. Therefore, we add these component offer prices, to which we refer to as *op*, to the existing regime model. These prices result from all requests for quotation in a TAC SCM game. These requests include requests for short-term delivery (e.g., 5 days) as well as long-term delivery (e.g., 50 days).

Regime Model Variables

As both regime identification and prediction are based on normalized sales prices, this section introduces a mathematical formulation of the normalized price np for product g , np_g , on day d , which is calculated as

$$np_g = \frac{\text{price}_g}{\text{asmCost}_g + \sum_{j=1}^{\text{numParts}_g} \text{nomPartCost}_{g,j}} , \quad (4)$$

using the product price price_g and the nominal manufacturing costs for each component j belonging to product g , $\text{nomPartCost}_{g,j}$, respectively.

The estimated normalized mean price can be volatile and lacks information on trends. Therefore, (exponential) smoothing can be applied, resulting in yesterday’s smoothed normalized minimum and maximum prices $\widetilde{np}_{d-1}^{\min}$ and $\widetilde{np}_{d-1}^{\max}$. Equations (5) through (7) show the smoothing process using a Brown linear exponential smoother (Brown, Meyer, and D’Esopo 1961), where α is a smoothing factor determined by a hill-climbing procedure which minimizes the variance of $\widetilde{np}_{d-1}^{\min}$.

$$\widetilde{np}_{d-1}^{\min'} = \alpha \cdot np_{d-1}^{\min} + (1 - \alpha) \cdot \widetilde{np}_{d-2}^{\min'} , \quad (5)$$

$$\widetilde{np}_{d-1}^{\min''} = \alpha \cdot \widetilde{np}_{d-1}^{\min'} + (1 - \alpha) \cdot \widetilde{np}_{d-2}^{\min''} , \quad (6)$$

$$\widetilde{np}_{d-1}^{\min} = 2 \cdot \widetilde{np}_{d-1}^{\min'} - \widetilde{np}_{d-1}^{\min''} . \quad (7)$$

Here, two price components are smoothed separately, after which they are combined. This way, changes in the mean and trend can be captured. Brown linear exponential smoothing is applied, since the trend as well as the mean vary over time. The calculation of $\widetilde{\text{np}}_{d-1}^{\max}$ is done by analogy with (5) through (7). Now, yesterday's exponentially smoothed normalized price on an arbitrary day d can be calculated by averaging yesterday's exponentially smoothed normalized minimum and maximum prices, i.e., $\widetilde{\text{np}}_{d-1} = 0.5 \cdot \widetilde{\text{np}}_{d-1}^{\min} + 0.5 \cdot \widetilde{\text{np}}_{d-1}^{\max}$.

We extend the regime model with the mean component offer price for product g , op_g , on day d , such that

$$\text{op}_g = \frac{\sum_{s=1}^{\text{numS}} \sum_{c=1}^{\text{numC}_g} \text{op}_{g,s,c}}{\text{numOp}_g}, \quad (8)$$

where numS refers to the number of suppliers, numC_g refers to the number of components for product g , and numOp_g represents the number of entries of the procurement variable. Thus, the mean component offer price is calculated by means of a counting process over all component prices. These prices are extracted from all requests for quotation related to a specific product send on an arbitrary day.

Because we would like the variable to include some information about other preceding days as well, so that it represents a trend instead of an event, we apply an exponential smoother to the variable. The smoothed value of yesterday's product-based component offer price, $\widetilde{\text{op}}_{d-1}$, is calculated as shown in (9), where β represents a smoothing factor and is determined using a hill-climbing procedure which minimizes the variance of $\widetilde{\text{op}}_{d-1}$:

$$\widetilde{\text{op}}_{d-1} = \beta \cdot \text{op}_{d-1} + (1 - \beta) \cdot \widetilde{\text{op}}_{d-2}. \quad (9)$$

Smoothing is done by taking a certain percentage (β) of yesterday's value of op . Then, the remaining percentage is taken of the previous (smoothed) value of variable op , i.e., the day before yesterday's value, after which both values are added. This is a less complex way of smoothing than applies for the normalized mean sales price, though it still smoothes out the variable's possible volatility.

Regime Identification

The existing regime model is based on a Gaussian Mixture Model (GMM) (Titterton, Smith, and Makov 1985) with a fixed number (N) of Gaussian components. A GMM is used, since it is able to approximate arbitrary density functions. Also, a GMM is a semi-parametric approach which allows for fast computing and uses less memory than other approaches (Ketter et al. 2009). In the current regime model, fixed means, μ_i , which are equally distributed, and variances, σ_i^2 , where i is used as an index to point to a component, $\forall i = 1, \dots, N$, are used. The fixed means and variances are chosen so that adjacent Gaussians are two standard deviations apart (Ketter et al. 2008), are used. This might lead to good results when fitting a model on one

dimension, but after adding a dimension to the model, fixed means and variances might prevent the GMM to reach a good fit. Therefore, we do not constrain the means and variances for now.

As is the case with the current model, we apply the Expectation-Maximization algorithm to determine the Gaussian components of the GMM and their prior probabilities, $P(\zeta_i)$. The Gaussian components are, unlike the components of the current model, based on both np and op . For now, the number of Gaussian components, N , is equal to 3, because this helps visualizing and explaining the main concepts of the model. We define the bivariate density of the sales and procurement offer prices as

$$p(\text{np} \cap \text{op}) = \sum_{i=1}^N P(\zeta_i) p(\text{np} \cap \text{op} | \zeta_i). \quad (10)$$

This density is equal to the sum of all Gaussian components $p(\text{np} \cap \text{op} | \zeta_i)$ multiplied by their prior probabilities $P(\zeta_i)$. We define a typical two-dimensional Gaussian component as

$$\begin{aligned} p(\text{np} \cap \text{op} | \zeta_i) &= p(\text{np} \cap \text{op} | \{ \mu_{\text{np}_i} \cap \mu_{\text{op}_i} \cap \sigma_{\text{np}_i} \cap \sigma_{\text{op}_i} \}) \\ &= A e^{-\left(\frac{(\text{np} - \mu_{\text{np}_i})^2}{2\sigma_{\text{np}_i}^2} + \frac{(\text{op} - \mu_{\text{op}_i})^2}{2\sigma_{\text{op}_i}^2} \right)}, \end{aligned} \quad (11)$$

where A is the amplitude of the Gaussian density, μ_{np_i} and μ_{op_i} are the means of the i -th Gaussian on the normalized mean price and mean offer price axes, and σ_{np_i} and σ_{op_i} are their respective standard deviations.

Figure 1 shows a two-dimensional GMM created using a training set² and the equations discussed above. For sake of illustration, the model contains three Gaussian components, and is trained with a maximum of fifteen hundred iterations on data on the high market segment. Experiments show that using less iterations does not guarantee a well fit model. Figures 1(a) and 1(c) show projections of the individual Gaussians and the density of the normalized mean sales price and mean offer price onto the axes of both variables, whereas the density is shown as a surface in Figure 1(b).

In order to find patterns in these probabilities, we need to calculate to which extent each of the probabilities is a member of each Gaussian component. The posterior probability for each component $P(\zeta_i | \text{np} \cap \text{op})$ follows from (10) after applying Bayes' rule:

$$\begin{aligned} P(\zeta_i | \text{np} \cap \text{op}) &= \frac{P(\zeta_i) p(\text{np} \cap \text{op} | \zeta_i)}{\sum_{j=1}^N P(\zeta_j) p(\text{np} \cap \text{op} | \zeta_j)}, \\ &\forall i = 1, \dots, N. \end{aligned} \quad (12)$$

Equation (12) applies for each Gaussian, and thus the vector of posterior probabilities for the two-dimensional Gaussian Mixture Model can be described

²Training set contains 2007 Semi-Finals games played on the SICS tac5 server (IDs: 9323–9327), 2007 Finals games played on the SICS tac3 server (IDs: 7308–7312), 2008 Semi-Finals games played on the UMN tac02 server (IDs: 763–768), and 2008 Finals games played on the UMN tac01 server (IDs: 794–799).




Figure 1: A two-dimensional GMM based on np and op , using three Gaussian components, where (a) and (c) show projections of the Gaussian components used in the model demonstrated in (b). Each individual Gaussian has its own density function, and combining these densities results into a single price density.

as $\eta(np \cap op) = [P(\zeta_1|np \cap op), \dots, P(\zeta_N|np \cap op)]$. For each combination of normalized mean prices and component offer prices, we can calculate $\eta(np \cap op)$ using the fitted Gaussian Mixture Model.

We need to find clusters within the posterior probabilities, as they can be linked to regimes, because they each describe certain conditions and characteristics. Clustering the posterior probabilities in M clusters is done using **K-Means** clustering. We tested different clustering algorithms, such as spectral clustering, which resulted in similar clusters. Clustering is done in fifteen replicates, using a maximum of one hundred iterations. Experiments show that this maximum allows the algorithm to converge nicely on our data set. The squared Euclidean distance measure is used to measure distances to the cluster centers for each data point. Figure 2 shows results of applying the **K-Means** clustering algorithm to the GMM we have fit to our data on high-range products with three clusters.

We link the cluster centers $P(\zeta|R_k)$ to regimes, but these clusters do not tell us which cluster represents which regime. Let us assume for now we know how to

assign the proper regime label to each cluster. Then we can rewrite $p(np \cap op|\zeta_i)$ by analogy with (10) in a form that shows the dependence of the normalized sales price and mean component offer price on the regime R_k .

$$p(np \cap op|R_k) = \sum_{i=1}^N P(\zeta_i|R_k) p(np \cap op|\zeta_i), \quad \forall k = 1, \dots, M. \quad (13)$$

In (13), $P(\zeta_i|R_k)$ refers to the N by M matrix resulting from the **K-Means** algorithm, and $p(np \cap op|\zeta_i)$ refers to the individual Gaussians. When applying Bayes' rule, we obtain the probability of regime R_k dependent on the sales and offer prices, as defined in (14).

$$P(R_k|np \cap op) = \frac{P(R_k) p(np \cap op|R_k)}{\sum_{j=1}^M P(R_j) p(np \cap op|R_j)}, \quad \forall k = 1, \dots, M. \quad (14)$$

Figure 3 shows a plot of the regime probabilities (given normalized sales price and component offer price)




Figure 2: Three identified clusters in the posterior probabilities, $P(\zeta_i|np \cap op)$, of a two-dimensional GMM.




Figure 3: Regime probabilities, $P(R_k|np \cap op)$, for products of the high segment.

for products of the low segment, resulting from a Gaussian Mixture Model and clustering its posterior probabilities in three clusters. We observe that under different conditions, different regimes are dominant, as different clusters have high probabilities for different combinations of component offer and sales order prices. Thus, each identified regime is dominant for certain combinations of both variables the model is based on.

Online, regime probabilities can be calculated (interpolated) for different combinations of normalized mean prices and procurement-side offer prices using the data that is displayed in Figure 3. Then, the dominant regime \widehat{R}_r on an arbitrary game day d is

$$\widehat{R}_r \text{ s.t. } r = \underset{1 \leq k \leq M}{\operatorname{argmax}} \widetilde{P}(\widehat{R}_k | \widetilde{\text{np}}_{d-1} \cap \widetilde{\text{op}}_{d-1}). \quad (15)$$

We conclude that in general, the regime identification still works similar to the existing approach. However, a dimension has been added to the Gaussian Mixture Model, causing differently structured probability densities as well as regime clusters. This requires reformulating the entire regime identification model.

Regime Prediction

We already introduced three techniques for regime prediction, each of which has its own characteristics and optimal time span to predict regimes for. This section continues with discussing these techniques. First, we will elaborate on the exponential smoother process, after which we will discuss Markov processes.

The exponential smoother regime prediction process is more reactive to the current market condition than any other method, because the exponential smoother process takes yesterday's information as input. This information is smoothed with information on preceding days to reduce volatility.

The prediction process calculates a trend, $\widetilde{\text{tr}}_{d-1}^{\min}$, in the minimum normalized mean sales price by using (5) and (6), and smoothing factor γ : $\widetilde{\text{tr}}_{d-1}^{\min} = \frac{\gamma}{1-\gamma} \cdot (\widetilde{\text{np}}_{d-1}^{\min'} - \widetilde{\text{np}}_{d-1}^{\min''})$. The exponentially smoothed maximum normalized trend, $\widetilde{\text{tr}}_{d-1}^{\max}$, is calculated in a similar way. Using the minimum and maximum trends, the mid-range trend of the sales price can be calculated by averaging both values, and thus $\widetilde{\text{tr}}_{d-1}^{\text{np}} = 0.5 \cdot \widetilde{\text{tr}}_{d-1}^{\min} + 0.5 \cdot \widetilde{\text{tr}}_{d-1}^{\max}$.

Using yesterday's value and the mid-range trend of sales prices, one can estimate the value of sales prices n days in the future as shown in (16), where h is the planning horizon:

$$\widetilde{\text{np}}_{d+n} = \widetilde{\text{np}}_{d-1} + (1+n) \cdot \widetilde{\text{tr}}_{d-1}^{\text{np}}, \quad \forall n = 0, \dots, h. \quad (16)$$

Now that we have defined a way to predict future values of np, we can also formulate a way to predict future values of op. This is also done using a trend, as the approach performs good in the current model and we prefer to keep things similar to the current approach. Also, it is beyond our scope to look into alternatives. However, the calculation is somewhat different from what

we have discussed for np, because of the fact that the procurement variable (i.e., component offer prices) represents a mean value and we do not have minimum and maximum values at hand. Also, different smoothing is applied to offer prices than to sales prices, which means the two Brown linear exponential smoothing components used for calculating the sales price trend are not available for our procurement variable.

We calculate the trend of $\widetilde{\text{op}}_{d-1}$, by computing the difference between yesterday's smoothed component offer price and the smoothed offer price of the day before yesterday, i.e., $\widetilde{\text{tr}}_{d-1}^{\text{op}} = \widetilde{\text{op}}_{d-1} - \widetilde{\text{op}}_{d-2}$. Then, future values for n days into the future up to planning horizon h are calculated similar to future values of $\widetilde{\text{np}}$, as shown in (17). Here, the calculated trend is added $1+n$ times to the last known value of the component offer prices. We express the future values of $\widetilde{\text{op}}$ mathematically as

$$\widetilde{\text{op}}_{d+n} = \widetilde{\text{op}}_{d-1} + (1+n) \cdot \widetilde{\text{tr}}_{d-1}^{\text{op}}, \quad \forall n = 0, \dots, h. \quad (17)$$

Using the future values for day $d+n$ of $\widetilde{\text{np}}$ and $\widetilde{\text{op}}$, the probability for each regime (given both prices) for n days in the future can be calculated similar to (14), where the density of the variables dependent on regime R_k is calculated using (13) by marginalizing over the individual Gaussians and cluster centers:

$$P(\widehat{R}_k | \widetilde{\text{np}}_{d+n} \cap \widetilde{\text{op}}_{d+n}) = \frac{p(\widetilde{\text{np}}_{d+n} \cap \widetilde{\text{op}}_{d+n} | \widehat{R}_k) P(R_k)}{\sum_{j=1}^M p(\widetilde{\text{np}}_{d+n} \cap \widetilde{\text{op}}_{d+n} | \widehat{R}_j) P(R_j)}, \quad \forall k = 1, \dots, M. \quad (18)$$

Making short-term and long-term regime predictions can be done using Markov prediction and correction-prediction processes (Isaacson and Madsen 1976). In contrast to the exponential smoother process where future prices are predicted, resulting indirectly in predictions of future regimes, regimes are predicted directly. Markov processes are less responsive to current market situations, as they make use of a history of events.

For short-term regime predictions, a Markov prediction process is used. This process is based on the last price measurement and on a Markov transition matrix $T(r_{d+n}|r_d)$. This matrix is created by means of a counting process on offline data and contains posterior probabilities of transitioning to regime r_{d+n} on day $d+n$ (i.e., n days into the future), given r_d , which is the current regime. Note that we introduce the symbol r for denoting regimes, to emphasize a focus shift, as we are not looking at the individual regimes in the way we were looking at them until now. Now, the regimes represent rows and columns in a Markov transition matrix and we use probability vectors combined with transition matrices, instead of single regime probabilities.

Ketter et al. distinguish between two types of Markov predictions: n -day prediction and repeated one-day prediction. The first type is an interval prediction, where

for each day n up to planning horizon h a Markov transition matrix is computed offline (per product, or at whatever level of detail the regime model is defined), whereas the second type only needs one Markov transition matrix (per product). This matrix is repeated n times up to planning horizon h .

The calculation of the n -day prediction for n days ahead is performed recursively as

$$\begin{aligned} \vec{P}(\hat{r}_{d+n}|\hat{\text{np}}_{d-1} \cap \hat{\text{op}}_{d-1}) = & \\ \sum_{r_{d+n}} \dots \sum_{r_{d-1}} \left\{ \vec{P}(\hat{r}_{d-1}|\hat{\text{np}}_{d-1} \cap \hat{\text{op}}_{d-1}) \cdot \right. & \\ \left. T_n(r_{d+n}|r_{d-1}) \right\}, & \\ \forall n = 0, \dots, h, & \end{aligned} \quad (19)$$

where the previous (identified or predicted) posterior regime probabilities dependent on the normalized mean sales prices and normalized mean component offer prices are multiplied with the applicable Markov transition matrix. The calculation of the repeated one-day prediction is done similarly, as shown in (20). However, only one transition matrix is used, $T_0(r_d|r_{d-1})$, which is repeated n times for predictions of n days in the future.

$$\begin{aligned} \vec{P}(\hat{r}_{d+n}|\hat{\text{np}}_{d-1} \cap \hat{\text{op}}_{d-1}) = & \\ \sum_{r_{d+n}} \dots \sum_{r_{d-1}} \left\{ \vec{P}(\hat{r}_{d-1}|\hat{\text{np}}_{d-1} \cap \hat{\text{op}}_{d-1}) \cdot \right. & \\ \left. \prod_{t=0}^n T_0(r_d|r_{d-1}) \right\}, & \\ \forall n = 0, \dots, h. & \end{aligned} \quad (20)$$

Long-term predictions are made using a Markov correction-prediction process. The latter process is almost similar to the Markov prediction process we already discussed. The difference is that the long-term prediction process is modeled based on the entire history of prices, instead of just the last price measurement. Hence, we need to extend the probability of regime \hat{r}_{d-1} dependent on $\hat{\text{np}}_{d-1}$ and $\hat{\text{op}}_{d-1}$ so that it incorporates the entire history of the values of $\hat{\text{np}}$ and $\hat{\text{op}}$. This correction is done by applying a recursive Bayesian update to the identified regime probabilities. Then, predictions for today are based on the corrected regime probabilities on day $d-1$, while predictions for n days in the future are done recursively. Equation (21) defines the n -day variant of the Markov correction-prediction process.

$$\begin{aligned} \vec{P}(\hat{r}_{d+n}|\{\hat{\text{np}}_1, \dots, \hat{\text{np}}_{d-1}\} \cap \{\hat{\text{op}}_1, \dots, \hat{\text{op}}_{d-1}\}) = & \\ \sum_{r_{d+n}} \dots \sum_{r_{d-1}} \left\{ \vec{P}(\hat{r}_{d-1}|\{\hat{\text{np}}_1, \dots, \hat{\text{np}}_{d-1}\} \cap \right. & \\ \left. \{\hat{\text{op}}_1, \dots, \hat{\text{op}}_{d-1}\}) \cdot T_n(r_{d+n}|r_{d-1}) \right\}, & \\ \forall n = 0, \dots, h. & \end{aligned} \quad (21)$$

The same principles apply to the calculation of the repeated one-day correction-prediction. Again, the difference is the usage of the Markov transition matrices.

$$\begin{aligned} \vec{P}(\hat{r}_{d+n}|\{\hat{\text{np}}_1, \dots, \hat{\text{np}}_{d-1}\} \cap \{\hat{\text{op}}_1, \dots, \hat{\text{op}}_{d-1}\}) = & \\ \sum_{r_{d+n}} \dots \sum_{r_{d-1}} \left\{ \vec{P}(\hat{r}_{d-1}|\{\hat{\text{np}}_1, \dots, \hat{\text{np}}_{d-1}\} \cap \right. & \\ \left. \{\hat{\text{op}}_1, \dots, \hat{\text{op}}_{d-1}\}) \cdot \prod_{t=0}^n T_0(r_d|r_{d-1}) \right\}, & \\ \forall n = 0, \dots, h. & \end{aligned} \quad (22)$$

Note that the Markov transition matrices which are being used in (21) and (22) ($T_n(r_{d+n}|r_{d-1}) \forall n = 0, \dots, h$ and $T_0(r_d|r_{d-1})$, respectively) are the same matrices as used in (19) and (20).

Performance Evaluation

Initial offline experiments show that a five-regime model is preferred over a three-regime model, as the regimes are identified and predicted more accurately. In these experiments, we train regime models using our training data, after which the performance is evaluated using a test set³.

Regime Identification Evaluation

We evaluate each model on several aspects. Correlations between identified dominant regime and economic regime identifiers, such as factory utilization and finished goods, are evaluated to test the feasibility of the clusters that have been found. Finally, the feasibility of the course of regime probabilities is evaluated.

The best performing regime model is configured with five regimes and ten Gaussian components, with which we continue our experiments. Figure 5 shows an example of the course of its identified regime probabilities through an arbitrary TAC SCM game of a typical agent in the mid-range product segment. Regimes are clearly dominant for a certain time and regimes do not switch too often, which is also visible in Figure 4. Here, the daily dominant regimes of the same game are displayed, together with the course of the normalized mean sales price and some other economic identifiers.

Regime labels are assigned to the cluster centers by means of correlation studies, of which the results are shown in Figure 6. In these studies, seventy-five hundred data points drawn from the training set are used to calculate the Pearson correlation (Fisher 1925) between the identified dominant regime and economic regime identifiers. This number of samples is large enough to ensure p -values below 0.01. Characteristics of the clusters are in line with the human interpretation of the

³Test set contains 2007 Semi-Finals games played on the SICS tac5 server (IDs: 9321, 9322, 9328), 2007 Finals games played on the SICS tac3 server (IDs: 7306, 7307, 7313), 2008 Semi-Finals games played on the UMN tac02 server (IDs: 761, 762, 769), and 2008 Finals games played on the UMN tac01 server (IDs: 792, 793, 800).




Figure 4: Overview of the dominant regimes during a TAC SCM game, together with the course of the normalized mean sales price and other economic identifiers, i.e., finished goods and factory utilization.

regime definitions. For example, during scarcity, there is a shortage of finished goods, and sales prices are high.

Regime Prediction Evaluation

We evaluate how the agent would predict the probabilities for five regimes with our ten-Gaussian GMM using our test set that contains historical data, and we compare these results to those of the current implementation of the regime model on the same test set. In our experiments, we evaluate three product segments, i.e., the low-, mid-, and high-range segment, to get a rough indication of the performance. This performance is measured in terms of the percentage of correctly predicted regimes and regime switch occurrences. We define the correct regime as the identified regime.

Regimes are predicted using a combination of the prediction methods we have discussed. Today's regime probabilities are predicted using an exponential smoother process. Short-term predictions (up to ten days into the future) are done using a Markov prediction process. We choose to use the n -day variant, since repeated one-day prediction tends to converge to a stationary distribution sooner. Finally, for long-term predictions up to twenty days into the future, we apply a Markov correction-prediction process. The upper bound (or planning horizon h) is set to twenty days, because production scheduling is set up every day for the next twenty days. This possibly leads to new or more accurate insights in future developments.




Figure 5: Course of identified regime probabilities over game time in an arbitrary TAC SCM game.




Figure 6: Correlation coefficients of five identified regime clusters, resulting from a two-dimensional Gaussian Mixture Model with ten individual Gaussians created based on mid-range product data.

Looking at the prediction performance of the selected regime model (compared to the performance of the current model), one can observe small improvements, as well as small deteriorations. This observation is supported by Table 2. Here, the accuracy measured in a percentage of correctly predicted regimes and regime switches (within plus or minus two days). The table shows the performance of the new model compared to the current model for three market segments (i.e., low-range, mid-range, and high-range products). The score of the best performing model is printed bold.

Correct	Segment	New model	Existing model
Regime	Low	46.43%	51.86%
	Mid	40.63%	52.93%
	High	40.48%	41.91%
Time	Low	48.68%	52.78%
	Mid	53.00%	43.44%
	High	50.93%	46.30%

Table 2: Prediction performances of a two-dimensional GMM with five clusters and ten individual Gaussians (new model) compared to the performance of a one-dimensional five-cluster GMM with sixteen individual Gaussians (existing model).

The differences between the scores of the existing model as presented in Table 2 and the results presented in (Ketter et al. 2006) can be caused by the fact that Ketter et al. only apply a Markov prediction process for each prediction. Furthermore, we experiment on 2007 and 2008 data, which contains different games than the ones used in (Ketter et al. 2006). This may result in market conditions which are harder to predict, because agents are getting more advanced and more competitive every year, which can cause other decisions to be made under the same conditions and thus games could have different characteristics.

On average, our model predicts regime switches more accurately than the current model. This indicates that the addition of procurement information does affect the prediction performance positively. However, regimes are predicted with a lower accuracy than the current model. Though overall, the differences between the performances are small, and therefore we conclude that the addition of procurement information in the proposed way does not affect the prediction performance greatly.

Conclusions and Future Work

We have added procurement information (component offer prices) to a sales-based regime model, which is used for predicting price density probabilities in a simulated supply chain. The regime model has been extended by adding a dimension to the Gaussian Mixture Model which is at the core of the economic regime model. After evaluating the performance of our model through experiments with the MinneTAC agent, which competes in the TAC SCM game for several years, we find that the new regime model has a similar overall

predictive performance as the existing model. Regime switches are predicted more accurately, whereas the prediction accuracy of dominant regimes is slightly worse.

However, by adding procurement information, we have enriched the model and we expect the new regime model to yield good results once implemented in the MinneTAC agent. The agent will be able to make decisions based on more information that is indicative of how market conditions are. Therefore, the agent will classify market conditions differently and more well-considered, which could lead to improvements. Also, our model seems to be as robust as the current model, and thus, maybe new types of decision making come within reach. Furthermore, we see opportunities for applications in the procurement market, which is worth further research.

For further research, we suggest to run tests against other competitors in the TAC SCM game to verify the model. We are currently looking into the use of time-delayed procurement information, as procurement information could be a leading indicator for the sales market. Subsequently, more or different procurement information can serve as a basis to the model. Not only using other data (e.g., offer quantities) or time-delayed data, but also applying other smoothing techniques to offer prices and normalizing data fall within the scope of the meaning of different procurement information.

References

- Andrews, J.; Benisch, M.; Sardinha, A.; and Sadeh, N. 2007. What Differentiates a Winning Agent: An Information Gain Based Analysis of TAC-SCM. In *AAAI Workshop on Trading Agent Design and Analysis (TADA'07)*.
- Becker, R.; Hurn, S.; and Pavlov, V. 2007. Modelling Spikes in Electricity Prices. *The Economic Record* 82(263):371–382.
- Brown, R. G.; Meyer, R. F.; and D’Esopo, D. A. 1961. The fundamental theorem of exponential smoothing. *Operations Research* 9(5):673–687.
- Collins, J.; Arunachalam, R.; Sadeh, N.; Eriksson, J.; Finne, N.; and Janson, S. 2005. The Supply Chain Management Game for the 2006 Trading Agent Competition. Technical Report CMU-ISRI-05-132, Carnegie Mellon University, Pittsburgh, PA.
- Collins, J.; Ketter, W.; and Gini, M. 2009. Flexible Decision Control in an Autonomous Trading Agent. *Electronic Commerce Research and Applications (ECRA)*.
- Fisher, R. A. 1925. *Statistical Methods for Research Workers*. Oliver & Boyd.
- Hamilton, J. D. 1989. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* 57(2):357–384.
- Isaacson, D. L., and Madsen, R. W. 1976. *Markov*

Chains – Theory and Applications. John Wiley & Sons.

Ketter, W.; Collins, J.; Gini, M.; Gupta, A.; and Schrater, P. 2006. Identifying and Forecasting Economic Regimes in TAC SCM. In Poutré, H. L.; Sadeh, N.; and Janson, S., eds., *AMEC and TADA 2005, LNAI 3937*. Springer Verlag Berlin Heidelberg. 113–125.

Ketter, W.; Collins, J.; Gini, M.; Gupta, A.; and Schrater, P. 2007. A predictive empirical model for pricing and resource allocation decisions. In *Proc. of 9th Int'l Conf. on Electronic Commerce*, 449–458.

Ketter, W.; Collins, J.; Gini, M.; Gupta, A.; and Schrater, P. 2008. Tactical and Strategic Sales Management for Intelligent Agents Guided By Economic Regimes. ERIM Working paper, ERS-2008-061-LIS.

Ketter, W.; Collins, J.; Gini, M.; Gupta, A.; and Schrater, P. 2009. Detecting and Forecasting Economic Regimes in Multi-Agent Automated Exchanges. *Decision Support Systems*.

Ketter, W. 2007. *Identification and Prediction of Economic Regimes to Guide Decision Making in Multi-Agent Marketplaces*. Ph.D. Dissertation, University of Minnesota, Twin-Cities, USA.

MacQueen, J. B. 1967. Some methods of classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability (MSP'67)*, 281–297.

Massey, C., and Wu, G. 2005. Detecting regime shifts: The causes of under- and overestimation. *Management Science* 51(6):932–947.

Mitchell, T. M. 1997. *Machine Learning*. McGraw-Hill. ISBN: 0071154671.

Mount, T. D.; Ning, Y.; and Cai, X. 2006. Predicting Price Spikes in Electricity Markets using a Regime-Switching model with Time-Varying Parameters. *Energy Economics* 28(1):62–80.

Osborn, D. R., and Sensier, M. 2002. The Prediction of Business Cycle Phases: Financial Variables and International Linkages. *National Institute Econ. Rev.* 182(1):96–105.

Titterton, D.; Smith, A.; and Makov, U. 1985. *Statistical Analysis of Finite Mixture Distributions*. New York: John Wiley and Sons.