

A Predictive Empirical Model for Pricing and Resource Allocation Decisions

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

John Collins, Maria Gini*,
Paul Schrater
Computer Science and Engr.
University of Minnesota
{jcollins, gini,
schrater}@cs.umn.edu

Alok Gupta
Carlson Sch. of Mgmt.
University of Minnesota
gupta@csom.umn.edu

ABSTRACT

We present a semi-parametric model that describes pricing behaviors in a market environment, and we show how that model can be used to guide resource allocation and pricing decisions in an autonomous trading agent. We validate our model by presenting experimental results obtained in the Trading Agent Competition for Supply Chain Management.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence; K.4.4 [Computers and Society]: E-commerce

General Terms

Algorithms, Economics, Theory

Keywords

Agent-mediated electronic commerce, machine learning, market forecasting, dynamic pricing

1. INTRODUCTION

Organizations seeking business advantage are increasingly looking to automated decision support systems. Eventually, we envision such systems evolving into software agents that can act rationally on behalf of their users in a variety of application areas. We are especially interested in systems that can support decision processes associated with product sales. If we can predict price trends in the markets we serve, we can allocate scarce inventory and production resources where they will earn the greatest profit. Once we know what we have available to sell, we would also like to find

*Partially supported from the National Science Foundation under award NSF/IIS-0414466.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICEC'07, August 19–22, 2007, Minneapolis, Minnesota, USA.
Copyright 2007 ACM 978-1-59593-700-1/07/0008 ...\$5.00.

the highest possible price points that will sell the product quantities we have available.

We present an approach whereby an autonomous agent is able to make tactical decisions, such as product pricing, as well as strategic decisions, such as product mix and production planning, in order to maximize its expected profit in an uncertain market. The agent predicts future market conditions and adapts its decisions on procurement, production, and sales accordingly.

We validate our approach using the Trading Agent Competition for Supply Chain Management [3] (TAC SCM). In a TAC SCM game, each of the competing agents plays the part of a manufacturer of personal computers. Agents compete with each other in a procurement market for computer components, and in a sales market for customers, as shown in Figure 1. A typical game runs for 220 simulated days over about an hour of real time. Each agent starts with no inventory and an empty bank account, and must borrow (and pay interest) to build up an initial parts inventory before it can begin assembling and shipping computers. The agent with the largest bank account at the end of the game is the winner.

Customers express demand each day by issuing a set of RFQs for finished computers. Each RFQ specifies the type of computer, a quantity, a due date, a reserve price, and a penalty for late delivery. Each agent may choose to bid on some or all of the day's RFQs. For each RFQ, the bid with the lowest price will be accepted, as long as that price is at or below the customer's reserve price. Once a bid is accepted, the agent is obligated to ship the requested products by the due date, or it must pay the stated penalty for each day the shipment is late. Agents do not see the bids of other agents, but they see the daily high and low prices for each computer type from the previous day, and aggregate market statistics are supplied to the agents every 20 days via a market report. Customer demand varies through the course of the game by a random walk.

Agents assemble computers of different types from parts, which must be purchased from suppliers, and manage inventories of parts and finished goods. When agents wish to procure parts, they issue RFQs to individual suppliers, and suppliers respond with bids that specify price and availability. Suppliers ship ordered parts on or after the due date (supplier capacity is variable). Supplier prices are based on their current uncommitted capacity.

The paper is organized as follows. After a short overview of existing methods, we introduce our method of economic

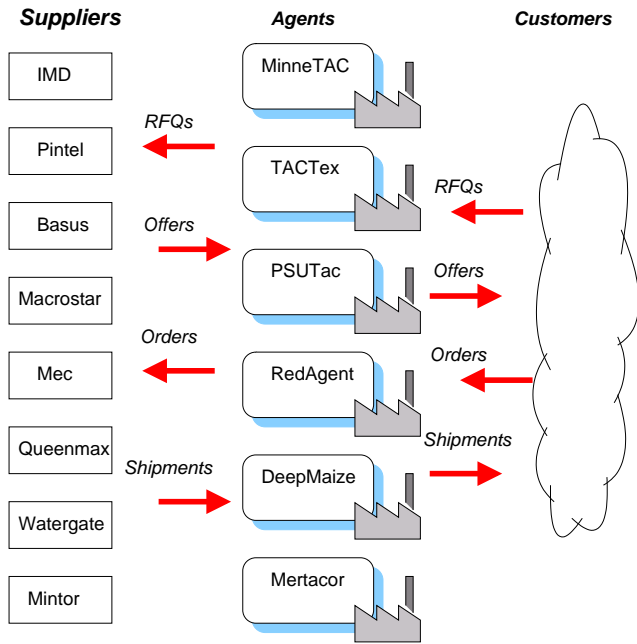


Figure 1: Schematic overview of a typical TAC SCM game scenario.

regimes for forecasting future market conditions and for modeling customer purchase behavior. In Section 4 we discuss an optimization approach to sales decisions, in particular for resource allocation and pricing. We present results of performance evaluation in Section 5. We conclude with ideas for future work.

2. EXISTING METHODS

The problem of predicting the probability of order in a sealed bid auction is commonly approached through statistical methods, such as those surveyed in [18]. These kinds of methods require large amounts of observed data on opponents bidding behavior and a static environment. TAC SCM, however, is a highly dynamic and uncertain environment and therefore nearly all agents in the TAC SCM competition use some dynamic way of modeling the probability of receiving an order.

Botticelli [1] uses a linear cumulative distribution function (CDF) to determine the relationship between offer price and order probability. TacTex [19] uses the lowest and highest offer price, which are provided for each product every day by the game server, and determines the probability of an order by linear interpolation.

RedAgent [12] uses an internal marketplace structure with competing bidders to set offer prices. PackaTAC [5] lets other agents set the price and tries to follow. The Jackaroo team [24] uses a game theoretic approach to set offer prices, using a variation of the Cournot game for modeling the product market. PSUTAC [22] employs an expert system for decision making, which allows them to express market strategies and knowledge in a human-understandable form. The SouthamptonSCM team [9] uses 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 [17], or PDA's on eBay [7]. The problem of allocating finite resources to producing a set of products in a way that maximizes some measure of utility is the well-known "product-mix" problem [10].

In [16] the authors demonstrate a method for predicting future customer demand in the TAC SCM environment, and use the predicted future demand to inform agent behavior. Their approach is specific to TAC SCM, since it depends on knowing the formula by which customer demand is computed. Note that customer demand is only one of the factors that characterizes the multi-dimensional regime parameter space we use.

Pindyck et al. [20] give a good overview of the science and art of building and using forecast models.

3. ECONOMIC REGIMES

Our goal is a model that can predict both future market prices and the probability that a particular bid price will result in an order. Economic theory suggests that economic environments tend to exhibit "dominant patterns" over time, such as scarcity, balanced situation, and oversupply. These patterns correspond to price ranges and trends. We call them *regimes*. We demonstrate a method for characterizing regimes using Gaussian Mixture Models (GMMs) and for learning them from market data. A complete description of this method is available in [13].

3.1 Regime model

Our approach begins with analyzing off-line data from past sales transactions in the market. We assume that historical data are available and that they are sufficiently representative of current market conditions. Information observable in real-time in the market is then used online to identify the current regime and to forecast regime transitions.

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

$$np_g = \frac{price_g}{nominalCost_g} \quad (1)$$

where $nominalCost_g$ is the "nominal" cost to source a unit of the good g . In TAC-SCM this is a fixed value for each good. An advantage of using normalized prices is that we can easily compare price patterns across different goods. In the following for simplicity we just use np .

Historical data are used to estimate the price density, $p(np)$, and to characterize regimes. We estimate the price density by fitting a Gaussian mixture model (GMM) [23] to historical normalized price data. We use a GMM since it is able to approximate arbitrary density functions. Since the GMM is a semi-parametric model it allows for fast computing using less memory than other approaches.

The results we present are obtained using a GMM with fixed means, μ_i , and fixed variances, σ_i^2 . The means are uniformly distributed and the variances are chosen so that adjacent Gaussians are two standard deviations apart. The reason for these choices is that they produce a set of Gaussians that works for all games. We use the Expectation-Maximization (EM) Algorithm [6] to determine the prior probability, $P(\zeta_i)$, of the Gaussians components of the GMM.

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 2.




Figure 2: Price density function modeled using 16 Gaussian components. Data are from 18 games from the semi-finals and finals of TAC SCM 2005

Given such a model for price density $p(\text{np})$, we apply Bayes' rule to determine the posterior probability of each Gaussian given a particular normalized price observation:

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

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

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

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

The intuitive idea of a regime as a recurrent economic condition is captured by discovering price distributions that recur across days. We define *regimes* by clustering the price distributions across days. This is done with the k-means algorithm. The resulting clusters correspond to frequently occurring price distributions with support on a contiguous range of np .

The center of each cluster is a probability vector that corresponds to regime $r = R_k$ for $k = 1, \dots, M$, where M is the number of regimes.

We selected the number of regimes a priori, after examining the data and looking at economic analysis of market

situations. In our experiments we found out that the number of regimes chosen does not significantly affect the results regarding price predictions.

Collecting the centers of the clusters into a matrix yields a conditional probability clustering matrix, which has N rows, one for each component of the GMM, and M columns, one for each regime.

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

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

The probability of regime R_k dependent on the normalized price np can then 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 historical data. Figure 3 depicts the regime probabilities for a sample market in TAC SCM. Each regime is clearly dominant over a range of normalized prices.




Figure 3: Learned regime probabilities $P(R_k|\text{np})$ over normalized price np , for a sample market after training.

In Figure 3 we distinguish five regimes, which we call extreme oversupply (R_1 or EO), oversupply (R_2 or O), balanced (R_3 or B), scarcity (R_4 or S), and extreme scarcity (R_5 or ES). We use five regimes instead of the usual three (oversupply, balanced, and scarcity) because this enables us to better isolate outlier situations.

Regimes EO and O represent a situation where there is an oversupply situation, which depresses prices. Regime B represents a balanced market situation, where most of the demand is satisfied. Regimes S and ES represent a situation where there is scarcity of goods in the market, which increases prices.

The intuition behind regimes is that prices communicate information about future expectations of the market. How-

ever, prices are of limited usefulness because the same price point can be achieved in a static mode (i.e., when prices are not changing), when prices are increasing, or when prices are decreasing. GMM and regimes capture information about price distributions.

3.2 Identifying regimes in real time

To do regime identification, we need to estimate the mean price of the goods sold. In TAC SCM, the agent receives every day a report that includes the minimum and maximum prices of goods sold the previous day, but not the quantities sold. The mid-range price, the price between the minimum and maximum, could be used as a coarse approximation of the mean price, but it is rather noisy.

To reduce the noise we smooth the minimum and the maximum prices using a Brown linear (double) exponential smoother [2] with $\alpha = 0.5$. In the following we show the smoothing of yesterday's minimum normalized price, np_{d-1}^{min} , and yesterday's maximum normalized price, np_{d-1}^{max} .

$$\widetilde{np}_{d-1}^{min'} = \alpha \cdot np_{d-1}^{min} + (1 - \alpha) \cdot \widetilde{np}_{d-2}^{min'} \quad (8)$$

$$\widetilde{np}_{d-1}^{min''} = \alpha \cdot \widetilde{np}_{d-1}^{min'} + (1 - \alpha) \cdot \widetilde{np}_{d-2}^{min''} \quad (9)$$

$$\widetilde{np}_{d-1}^{min} = 2 \cdot \widetilde{np}_{d-1}^{min'} - \widetilde{np}_{d-1}^{min''} \quad (10)$$

Since we only have the minimum and maximum prices from the previous day but not the mean, we model \widetilde{np}_{d-1} as follows:

$$\widetilde{np}_{d-1} = \frac{\widetilde{np}_{d-1}^{min} + \widetilde{np}_{d-1}^{max}}{2} \quad (11)$$

where $\widetilde{np}_{d-1}^{min}$ is the exponentially smoothed minimum normalized price and $\widetilde{np}_{d-1}^{max}$ is the exponentially smoothed maximum normalized price from the previous day. This results in a better approximation of the mean price than smoothing only the mid-range price from the previous day. Figure 4 shows that the smoothed mid-range price, \widetilde{np} , is closer to the mean price.

During the game, the agent estimates the current regime on day d by calculating the mean of the smoothed minimum and maximum prices \widetilde{np}_{d-1} for the previous day (recall that the agent receives each day the price observations for the previous day) and by selecting the regime which has the highest probability

$$\hat{R}_{max_1} \text{ s.t. } max_1 = \underset{1 \leq k \leq M}{\operatorname{argmax}} \vec{P}(\hat{R}_k | \widetilde{np}_{d-1}). \quad (12)$$

Figure 5 shows how to use the smoothed mid-range price, \widetilde{np} , to identify the corresponding regime probabilities online during the 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 correlation analysis of the market parameters to regimes, more details on regime identification, and other regime evaluation measures have been reported in [14, 15, 13].

3.3 Predicting future regimes

Since the behavior of an agent should depend not just on the current regime but also on expected future regimes, the agent needs to predict future regimes. We model the




Figure 4: Minimum, maximum, mean, mid-range, and smoothed mid-range daily normalized prices of computers sold every day, as reported during the game every day for the low market segment in the game 3717@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.




Figure 5: Regime probabilities computed online every day for the low market segment in game 3717@tac3.

prediction of future regimes for tactical decision making as a Markov prediction process [15]. The prediction is based on the last price measurement.

We construct a Markov transition matrix, $\mathbf{T}(r_{d+n}|r_d)$, offline by a counting process over past games. This matrix represents the posterior probability of transitioning on day $d+n$ to regime r_{d+n} given the current regime on day d , r_d . The prediction over n days is done by multiplying n times the one day prediction matrix, $\mathbf{T}_1(r_{d+1}|r_d)$.

The prediction of the posterior distribution of regimes n days into the future, $\vec{P}(\hat{r}_{d+n}|\widetilde{np}_{d-1})$, is done recursively as

follows:

$$\begin{aligned} \bar{P}(\hat{r}_{d+h}|\widehat{\text{np}}_{d-1}) = \\ \sum_{r_{d+h}} \cdots \sum_{r_{d-1}} \left\{ \bar{P}(\hat{r}_{d-1}|\widehat{\text{np}}_{d-1}) \cdot \mathbf{T}_1^{\text{h}+1}(r_d|r_{d-1}) \right\}, \end{aligned} \quad (13)$$

where

$$\mathbf{T}_1^{\text{h}+1}(r_d|r_{d-1}) = \prod_{n=0}^h \mathbf{T}_1(r_d|r_{d-1}) \quad (14)$$

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.

An alternative to the Markov prediction process is based on exponentially smoothed price predictions. In Equations 8 to 11 we have shown how to obtain an estimate of the smoothed mid-range price from the previous day, $\widehat{\text{np}}_{d-1}$. As a first step in predicting prices we calculate the price trend, tr_{d-1} , using Equation 8 and Equation 9, as:

$$\tilde{tr}_{d-1}^{\text{min}} = \frac{\alpha}{1-\alpha} \cdot (\widehat{\text{np}}_{d-1}^{\text{min}'} - \widehat{\text{np}}_{d-1}^{\text{min}''}) \quad (15)$$

Since we estimate the smoothed mid-range price as "the mean of the smoothed daily minimum and smoother maximum price (see Equation 11), we calculate the trend in a similar way as:

$$\tilde{tr}_{d-1} = \frac{\tilde{tr}_{d-1}^{\text{min}} + \tilde{tr}_{d-1}^{\text{max}}}{2} \quad (16)$$

where $\tilde{tr}_{d-1}^{\text{min}}$ is the exponentially smoothed minimum normalized trend and $\tilde{tr}_{d-1}^{\text{max}}$ is the exponentially smoothed maximum normalized trend from the previous day. Using the trend and yesterday's price estimate (see Equation 11) we estimate today's and the future daily smoothed prices as:

$$\widehat{\text{np}}_{d+n} = \widehat{\text{np}}_{d-1} + (1+n) \cdot \tilde{tr}_{d-1}, \quad \forall n = 1, \dots, h \quad (17)$$

We then obtain the density of the normalized price, $\widehat{\text{np}}_{d+n}$, dependent on the predicted regime \hat{R}_k :

$$p(\widehat{\text{np}}_{d+n}|\hat{R}_k) = \sum_{i=1}^N p(\widehat{\text{np}}_{d+n}|\zeta_i) P(\zeta_i|R_k). \quad (18)$$

The predicted probability of regime R_k dependent on the predicted exponentially smoothed normalized price n days into the future $\widehat{\text{np}}_{d+n}$ can be computed using the Bayes rule as:

$$\begin{aligned} P(\hat{R}_k|\widehat{\text{np}}_{d+n}) = \\ \frac{p(\widehat{\text{np}}_{d+n}|\hat{R}_k) P(R_k)}{\sum_{k=1}^M p(\widehat{\text{np}}_{d+n}|\hat{R}_k) P(R_k)} \quad \forall k = 1, \dots, M. \end{aligned} \quad (19)$$

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.

3.4 Predicting price distribution and trend

We now describe a method to predict the price trend based on regime prediction. An agent can use this information for guiding its future procurement, production, and pricing decisions.

Equation 20 describes how to compute a price prediction distribution based on a given predicted regime distribution.

As before M represents the number of regimes and N the number of Gaussians used in the GMM (see Equation 2). A point on the distribution, is given by

$$\begin{aligned} p(\widehat{\text{np}}_{d+n}|\widehat{\text{np}}_{d-1}) \\ = \sum_{i=1}^M p(\text{np}|R_i) P(\hat{R}_{i,d+n}|\widehat{\text{np}}_{d-1}) \\ = \sum_{j=1}^N \sum_{i=1}^M \underbrace{P(\zeta_j|R_i) P(\hat{R}_{i,d+n}|\widehat{\text{np}}_{d-1})}_{P(\zeta_j,d+n)} p(\text{np}|\zeta_j) \\ = \sum_{j=1}^N P(\zeta_j,d+n) p(\text{np}|\zeta_j), \quad \forall n = 1, \dots, h \end{aligned} \quad (20)$$

where $P(\hat{R}_{i,d+n}|\widehat{\text{np}}_{d-1})$ is one element of the predicted regime probability vector given by the Markov prediction or by the exponentially smoothed predicted regimes (see Equation 19).

To obtain a predicted price distribution we sample Equation 20 for every day over the planning horizon h with values for np between 0 and 1.25¹. After sampling the mixture distribution over the set of np values, the distribution is normalized to sum to one.

Figure 6 shows the forecast price density, based on a 1-day trained Markov matrix, for game 3717@tac3, for 20 days starting at day 113. The dashed curve represents the price density for the first forecasted day, the thick solid line shows the price density for the last forecasted day, and the thin solid curves show the forecast for the intermediate days. As expected, the predicted price density broadens as we forecast further into the future, reflecting a decreasing certainty in the prediction.

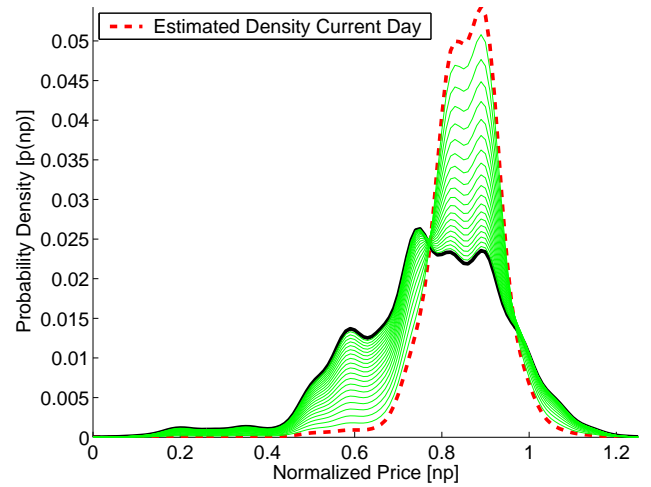


Figure 6: Predicted price density using the trained Markov matrix for game 3717@tac3 from day 113 until day 133.

We can compare the actual price trends with our predictions. Figure 7 shows the real mean price observations along with Markov and Exponential Smoother forecasts from day 113. To illustrate our predicted price trends, we show the

¹1.25 is the maximum normalized price that customers in TAC SCM are ever willing to pay.

10%, 50%, and 90% percentile of the predicted price distribution, which are interpolated from the discretized cumulative distribution. All the curves in the figure represent a relative price trend – to better compare the different predictors which each other graphically, we subtracted from each forecasted value the first predicted value, so that they all start at zero.

The exponential smoother prediction in this example is rather poor, since the smoother is myopic — it puts excess weight on recently observed prices. Figure 4 shows that prices before day 113 were increasing, but the exponential smoother predictor takes only the recent slope and extrapolates it into the future. On the contrary, our Markov predictor does a better job at learning during training how long, depending on the preconditions, a particular regime is active, and hence is a better predictor.

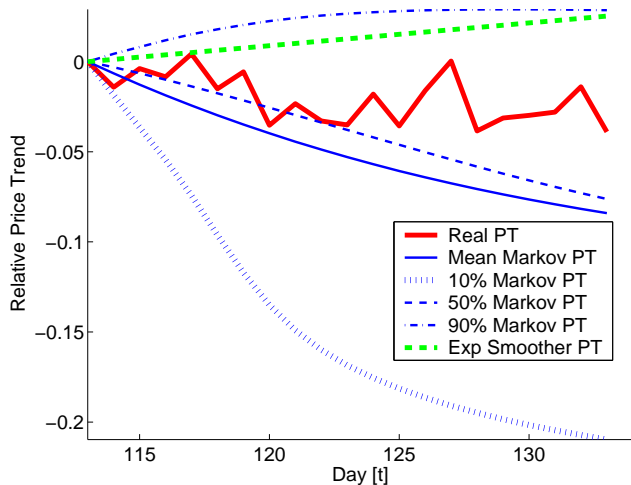


Figure 7: Predicted prices using the trained Markov matrix and exponential smoother for game 3717@tac3 from day 113 until day 133. The solid curve is the real mean price and the dashed and dotted curves are predicted price trends based on the 10%, 50% and 90% percentile of the predicted price density.

Given the estimate of the price density, the order probability function $P(\text{order}|\text{np})$ can be estimated as $1 - CDF(\text{np})$, where $CDF(\text{np})$ is the cumulative distribution function of np , the integral of the price density.

4. SALES DECISIONS

Given a model for future market prices, we would like to be able to maximize the profit our agent can expect to earn over some reasonable period in the future. We do this first by allocating limited resources over time and over alternate product mix options, and second by setting prices in a way that maximizes the value of the goods we wish to sell during each sales cycle.

Because procurement is at least partly driven by projected inventories and by predicted customer demand, sales activity also influences procurement. But procurement typically operates over a longer time horizon, and sales must be focused on getting the highest possible prices for the resources it has available. If inventory is sold out during a period when

prices are low, then there may be nothing available to sell when prices recover.

4.1 Resource allocation

Resource allocation decisions can be informed by experience in the past, and by observations in the present. The economic regime model encapsulates past experience and observes current market price data, giving us current and expected market prices. Other observable data includes the following:

- \mathcal{C} is the set of all available component types, and \mathcal{G} is the set of all goods (product types) that can be built and sold. Each good is made up of a set \mathcal{C}_g of components. This means that in turn, each component c is a part of some set of products \mathcal{G}_c .
- On each day d , customer demand is represented by a set \mathcal{R}_d of customer RFQs received by the agent. Each RFQ $r \in \mathcal{R}_d$ specifies a product type g_r , a lead time of i_r days, a quantity q_r , and a reserve price ρ_r . Reserve price is uniformly distributed between ρ^{\min} and ρ^{\max} . Details and semantics are given in [3].
- Customer demand is projected into the future over some planning horizon h . We model customer demand following the Bayesian inversion method described in [16].
- At the beginning of any given day d , the agent has an inventory of raw materials consisting of $I_{d,c}$ for each component type $c \in \mathcal{C}$, and an inventory of finished goods consisting of $I_{d,g}$ for each type of good $g \in \mathcal{G}$.

From this data, we would like to find a way to set prices and make offers to customers that maximize the agent's overall profits. On any given day d , the total demand $D_{d,g}$ for a given good g among $\mathcal{R}_d^{\text{cust}}$ is the total of the requested quantities among requests for good g , given by

$$D_{d,g} = \sum_{r \in \mathcal{R}_{d,g}^{\text{cust}}} q_r \quad (21)$$

We assume that the price $\text{price}_{d,g} = f(D_{d,g}, A_{d,g})$ sustainable by the market for a given good g on a particular day d is a function of the demand $D_{d,g}$ and the quantity of goods the agent wishes to sell, the *allocation* or sales quota $A_{d,g}$ for good g on day d .

The profit per unit for good g to be sold on day d at a price $\text{price}_{d,g}$ is given by

$$\Phi_{d,g} = \text{discount}(d)(\text{price}_{d,g} - \text{cost}(\mathcal{C}_g)) \quad (22)$$

We include the discount term as a rough approximation of inventory holding cost. It can also be used to encourage early selling, as a hedge against the uncertainty of the game.

For any given day d , there is an unsold inventory I'_g of good g , and an expected uncommitted inventory $I'_{d,c}$ of parts of type c . This includes parts in current inventory, and parts that are expected to be delivered by day d , and excludes parts that are allocated to produce goods for outstanding customer orders.

The *effective demand* function $D_{d,g}^{\text{eff}} = f(D_{d,g}, \text{price}_{d,g})$ for our goods will be some function of the prices $\text{price}_{d,g}$ we wish to charge. In the TAC SCM environment, there is a linear distribution of reserve prices among customer RFQs. The effective demand, then, is the portion of total demand

with reserve prices $\rho \geq price_{d,g}$ at or above the price we want to sell at:

$$D_{d,g}^{eff} = \frac{\rho_g^{\max} - price_{d,g}}{\rho_g^{\max} - \rho_g^{\min}} D_{d,g} \quad (23)$$

where ρ_g^{\max} is the maximum reserve price for good g . This assumes that actual demand is uniformly distributed across the range of reserve prices, which is only approximately correct.

The total profit Φ over a planning horizon of h days for the set of goods \mathcal{G} is then

$$\Phi = \sum_{d=0}^h \sum_{g \in \mathcal{G}} \Phi_{d,g} A_{d,g} \quad (24)$$

This is what we wish to maximize, by computing values for $A_{d,g}$, subject to the following constraints:

1. We can't sell more of any product than the effective demand at the price we wish to charge:

$$\forall d, \forall g, A_{d,g} < D_{d,g}^{eff} \quad (25)$$

2. For any given period of time from now until the planning horizon h , we can sell goods that we have in inventory, and goods for which we have the necessary parts in inventory.

$$\forall m \in 1..h, \forall c \in \mathcal{C}, \sum_{d=1}^m \sum_{g \in \mathcal{G}_c} A_{d,g} \leq I'_{m,c} + \sum_{g \in \mathcal{G}_c} I'_g \quad (26)$$

Note that this constraint limits commitments of the *sets* of goods that share a given component. If we don't carry any uncommitted finished goods inventory, in other words if $\forall g \in \mathcal{G}_c, I'_g = 0$, then this expresses the inventory constraints. Otherwise, imbalances in the finished-goods inventories of individual goods sharing a component could lead to overcommitment. This is easy to see if for some component c , $I'_{m,c} = 0$. Then the sum of individual product inventories constrain the whole set of products. In this case, it is also necessary to constrain every subset of product types that can share some component. This requires that we replace Equation 26 with

$$\forall m \in 1..h, \forall c \in \mathcal{C}, \forall \mathcal{G}'_c \subseteq \mathcal{G}_c, \sum_{d=1}^m \sum_{g \in \mathcal{G}'_c} A_{d,g} \leq I'_{m,c} + \sum_{g \in \mathcal{G}'_c} I'_g \quad (27)$$

3. The agent's factory has limited daily capacity F . If each unit of good g requires y_g production cycles, then

$$\forall m \in 1..h, \sum_{g \in \mathcal{G}} y_g \left(\sum_{d=1}^m A_{d,g} - I'_g \right) \leq mF - F_m^{commit} \quad (28)$$

where F_m^{commit} is the factory capacity that is committed to manufacture all outstanding customer orders that are due on or before day m and are not satisfiable by existing finished goods inventory.

The outcome of our objective function (Eq. 24) is daily sales quotas $A_{d,g}$ for each good. The next step is to set

prices so that we sell what we intend to sell, in a competitive market.

In Section 3.4, we described a model for estimating the probability of a customer placing an order as a function of price $P(order|price)$. But the quantity we sell is just the effective demand multiplied by the probability of order at the price we set. So to sell our sales quota, we need

$$A_{d,g} = P(order|price_{d,g}) D_{d,g}^{eff} \quad (29)$$

In the TAC SCM environment, with its linear distribution of reserve prices, this gives

$$A_{d,g} = P(order|price_{d,g}) \frac{\rho_g^{\max} - price_{d,g}}{\rho_g^{\max} - \rho_g^{\min}} D_{d,g} \quad (30)$$

which is quadratic in $price_{d,g}$, assuming that $P(order|price_{d,g})$ is linear (it is not). When we combine Equation 24 with Equations 22 and 30, we have an expression that is at least cubic in $price_{d,g}$.

Because the formula for sales quota allocations above is probably unsolvable given the time constraints of the TAC SCM simulation environment, and because we cannot in general assume that effective demand and order probability are linear functions, there is a need for heuristics and simplifications.

An obvious simplification is to assume that the partial derivative of the order probability function with respect to price is very steep. This is equivalent to saying that (most) sales occur at a "market clearing price," or alternatively that the probability of order is much more sensitive to price than is profit. Then the per-unit profit and the effective demand can be computed separately, by substituting an estimated clearing price $price_{d,g}^{est}$ for the actual sales price into Eq. 22. We will explore a way to compute $price_{d,g}^{est}$ in the next section, Eq. 35. Clearly this method depends on being able to project an estimate of $price_{d,g}^{est}$ over some period into the future.

4.2 Pricing

Once the strategic sales process has determined daily sales quotas, the next step is to set prices for our goods that will yield maximum profit. This amounts to finding, for each good, the value for $price_{d,g}$ that satisfies the relation

$$\frac{A_{d,g}}{D_{d,g}^{eff}} = P(order|price_{d,g}) \quad (31)$$

which is a simple rearrangement of Eq. 29.

This could be solved analytically or numerically, assuming we have reasonable functions for $D_{d,g}^{eff}$ and $P(order|price_{d,g})$. In general, however, one or both of these functions are likely to be empirically-derived histograms. Under the previous assumption of most sales occurring close to a market clearing price, we can approximate $D_{d,g}^{eff}$ using $price_{d,g}^{est}$, reducing the computation to finding the value of $price_{d,g}$ that satisfies

$$\frac{A_{d,g}}{D_{d,g}^{eff}(price_{d,g}^{est})} = P(order|price_{d,g}) \quad (32)$$

When prices are set in this way, the resulting customer orders provide an additional signal from the market that can be used to refine our estimate of $price_{d,g}^{est}$. If $O_{d,g}$ is the number of orders placed for good g on day d (as a result of offers made on day $d-1$), then a refined estimate of the actual market prices on that day $price_{d-1,g}^{act}$ can be found by finding

an adjusted probability distribution $P^{adj}(order|price_{d-1,g})$ such that

$$\frac{O_{d,g}}{D_{d-1,g}^{eff}} = P^{adj}(order|price_{d-1,g}) \quad (33)$$

and computing an estimated actual price $price_{d-1,g}^{act}$ such that

$$\frac{A_{d-1,g}}{D_{d-1,g}^{eff}} = P^{adj}(order|price_{d-1,g}^{act}) \quad (34)$$

See Fig. 8 for a graphical visualization of this relationship. For simplicity, we illustrate an approximate adjustment made by shifting the location of the probability curve along the price axis without changing its shape.




Figure 8: Estimating actual market price, given order volume O and an estimate of the order probability function P .

The resulting estimate $price_{d-1,g}^{act}$ is subject to the randomness of the market, and therefore we use an exponentially smoothed offset to produce a refined value of $price_{d,g}^{est}$ each day, as

$$price_{d,g}^{est} = price_{d,g}^{pred} + \delta_{d,g} \quad (35)$$

where $price_{d,g}^{pred}$ is the *predicted* market price for product g (see Sect. 3.4), and $\delta_{d,g}$ is updated daily as

$$\delta_{d,g} = \alpha \delta_{d-1,g} + (1 - \alpha)(price_{d-1,g}^{act} - price_{d-1,g}) \quad (36)$$

for some appropriate value of α .

5. PERFORMANCE EVALUATION

There are two ways to evaluate the power of our approach. The first is off-line, running the model on saved data and comparing its estimates and predictions to the actual data. The second is online, integrating the model into a working agent, and letting the agent play against other competitive agents. Here we describe a set of online experiments.

We implemented both the Markov prediction (MP) and exponential smoother (ES) prediction methods in a version of our MinneTAC agent [4], and ran multiple games against a set of well-known competitors.

The agents used in our experiments were obtained from the TAC SCM agent repository². In addition to MinneTAC, we selected four other finalists from the 2006 competition, and an agent from the 2005 competition. The agents are:

1. MinneTAC – University of Minnesota
2. DeepMaize – University of Michigan
3. Maxon – Xonar, Inc.
4. PhantAgent – Politechnica University of Bucharest
5. RationalSCM – Australian National University; competed in 2005.
6. TacTex – University of Texas; winner TAC SCM 2006

For our experiments we use a controlled server [21] to run N_G games, each with a different pseudo-random sequence, with MinneTAC and the five other agents, and then run N_G games with the *same market factors* (the same set of N_G pseudo-random sequences) with a modified MinneTAC and the same set of competing agents. In other words, all the pseudo-random sequences, as well as the set of agents competing with our test agent, from the first set of N_G games are repeated in the second set of N_G games. For our tests, $N_G = 23$. This method removes the profit variability due to the agents seeing different market conditions, and at the same time it removes the possibility of being hindered by unwanted interactions that can occur when multiple copies of agents under test are run against each other.

We use three different versions of our MinneTAC agent, each using different models for strategic decisions (price and price trend prediction) and for tactical decisions (order probability calculation).

For strategic decisions we used two different price prediction methods. The first is a price-follower method (an exponential smoother predicts future prices, without using a regime model), while the second uses regimes with Markov prediction as described in Section 3.4 (called “Regime-M” in Table 1).

For tactical decisions we used two methods to calculate the order probability. The first is a simple linear interpolation between the smoothed minimum and maximum prices, the second uses the regime model and makes predictions using the exponential smoother (called “Regime-E” in Table 1).

We used the Wilcoxon signed rank test [8, 11] to assess statistical significance among these three experiments. This is a *non-parametric* test of the difference between the medians of two samples that does not require the samples to come from normal (or even the same) distribution. The test is used to determine whether the median of a symmetric population is 0. First, the data are ranked without regard to sign. Second, the signs of the original observations are attached to their corresponding ranks. Finally, the one sample z statistic (mean / standard error of the mean) is calculated from the signed ranks.

Table 2 shows the results of the Wilcoxon test on our three experiment setups. The null hypothesis can be rejected if the medians from the two different samples are different. p is the probability of observing a result equally or more extreme than the one using the data (from both samples) if the null

²<http://www.sics.se/tac/showagents.php>

Test #	1: Exp 3 - Exp 1			2: Exp 3 - Exp 2			3: Exp 2 - Exp 1		
$\alpha = 0.05$	p	h	$srank$	p	h	$srank$	p	h	$srank$
All	0.0138	1	57	0.1137	0	86	0.0727	0	79
Positive	0.0054	1	13	0.2769	0	40	0.0256	1	21
Negative	0.4258	0	15	0.4258	0	15	0.9102	0	21

Table 2: Wilcoxon signed rank test of equality of medians. The tests were performed at a significance level of $\alpha = 0.05$ based on 23 data points. p represents the p-value, h is the result of the hypothesis test, $srank$ gives the value of the signed rank statistic.

Experiment	1	2	3
Strategic:	Follower	Regime-M	Regime-M
Tactical:	Linear	Linear	Regime-E
Agent	Mean Profit/Std. Dev. (in \$M)		
TacTex	8.75/5.68	8.87/5.60	9.21/5.39
DeepMaize	8.84/4.63	8.71/4.85	8.32/4.18
PhantAgent	8.05/5.42	7.99/5.38	8.17/5.44
Maxon	4.24/4.52	3.77/4.29	4.02/4.18
MinneTAC	1.35/3.70	1.81/4.02	2.12/3.76
Rational	0.74/4.91	0.67/4.69	1.31/4.53

Table 1: Experimental results with repeated market conditions and three variations of MinneTAC for order probability, price and price trend predictions. Mean profit and standard deviation results are based on 23 games. Regime-M uses the regime model with Markov prediction process, and Regime-E uses the regime model with exponential smoother lookup process.

hypothesis is true. If p is near zero, this casts doubt on the null hypothesis. The field “ $srank$ ” contains the value of the signed rank statistic.

We performed the tests on (1) the set of all the games, (2) only the positive profit games, and (3) only the negative profit games. As a result of these tests we observe that the results of experiments (3) and (2) are not significantly different, and we cannot reject the null hypothesis. On the other hand we find significant differences between the outcome of experiments (3) and (1).

We are able to reject the null hypothesis of equal median for the set of all games and the set of all positive games, but not for the set of negative games. The most likely reason why we are not able to reject the null hypothesis of equal median for the set of negative games is that in negative games an agent is more concerned with controlling cost than making profit and so the differences between the configuration is less apparent. We are also able to show significance between experiments (2) and (1) for the set of all positive games.

Although we could not reject the null hypothesis for the set of all games at $\alpha = 0.05$ significance level with a p-value of 0.0726, we are able to reject the null hypothesis for the set of all games at $\alpha = 0.1$ significance level. It is possible that the test would show significance with a larger sample size. We have shown statistical significance between the original configuration and the regime/regime configuration. Our results also suggest that the regime/regime configuration performs better than the linear/regime configuration, although more data would be needed to show that conclusively.

6. CONCLUSIONS

We have presented an approach for strategic and tactical decision making in a competitive sales environment, based on predicting market prices and price trends, optimizing product mix and resource allocation, and estimating the probability of receiving an order for a given offer price.

We have demonstrated the effectiveness of our approach by learning from games played in multiple TAC SCM games and by using the learned knowledge to predict prices and order probability during new games played with different agents. We have shown that using these methods results in significantly improved performance over a configuration that uses simpler methods, specifically a price-following predictor and a linear approximation of order probability. This improvement was achieved by modifying only the sales decision process, without any change to the decision processes for procurement, inventory management, or production scheduling.

We have implemented the regime identification and prediction methods in our TAC SCM agent and integrated it into the overall decision making process. Currently we are using regime predictions for tactical and strategic decision making in the sales component of our agent. Ultimately, we plan to combine probability information supplied by our method with information about possible consequences of actions to optimize the overall decision making, including inventory management, procurement, and production scheduling.

7. REFERENCES

- [1] M. Benisch, A. Greenwald, I. Grypari, R. Lederman, V. Naroditskiy, and M. Tschantz. Botticelli: A supply chain management agent designed to optimize under uncertainty. *ACM Trans. on Comp. Logic*, 4(3):29–37, 2004.
- [2] R. G. Brown, R. F. Meyer, and D. A. D’Esopo. The fundamental theorem of exponential smoothing. *Operations Research*, 9(5):673–687, 1961.
- [3] J. Collins, R. Arunachalam, N. Sadeh, J. Ericsson, N. Finne, and S. Janson. The supply chain management game for the 2006 trading agent competition. Technical Report CMU-ISRI-05-132, Carnegie Mellon University, Pittsburgh, PA 15213, November 2005.
- [4] J. Collins, W. Ketter, M. Gini, and A. Agovic. Software architecture of the MinneTAC supply-chain trading agent. Technical Report 07-006, University of Minnesota, Department of Computer Science and Engineering, Minneapolis, MN, February 2007.
- [5] E. Dahlgren and P. Wurman. PackaTAC: A

- conservative trading agent. *SIGecom Exchanges*, 4(3):33–40, 2004.
- [6] 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.
- [7] R. 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.
- [8] J. D. Gibbons. Nonparametric statistical inference. *Technometrics*, 28(3):275, 1986.
- [9] M. He, A. Rogers, E. David, and N. 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, Aug. 2005.
- [10] F. S. Hillier and G. J. Lieberman. *Introduction to Operations Research*. McGraw-Hill, 1990.
- [11] M. Hollander and D. A. Wolfe. Nonparametric statistical methods. *Journal of the American Statistical Association*, 95(449):333, 2000.
- [12] P. W. Keller, F.-O. Duguay, and D. Precup. Redagent - winner of the TAC SCM 2003. *SIGecom Exchanges*, 4(3):1–8, 2004.
- [13] W. 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.
- [14] W. Ketter, J. Collins, M. Gini, A. Gupta, and P. Schrater. A Computational Approach to Predicting Economic Regimes in Automated Exchanges. In *Proc. of the Fifteenth Annual Workshop on Information Technologies and Systems*, pages 147–152, Las Vegas, Nevada, USA, December 2005.
- [15] W. Ketter, J. Collins, M. Gini, A. Gupta, and P. Schrater. Detecting and forecasting economic regimes in automated exchanges. Technical Report 07-008, University of Minnesota, Dept of Computer Science and Engineering, Minneapolis, MN 55455, 2007.
- [16] C. Kiekintveld, M. P. Wellman, S. Singh, J. Estelle, Y. Vorobeychik, V. Soni, and M. Rudary. Distributed feedback control for decision making on supply chains. In *Fourteenth International Conference on Automated Planning and Scheduling*, Whistler, BC, Canada, June 2004. AAAI, AAAI Press.
- [17] R. 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, Jan. 2003.
- [18] V. Papaioannou and N. Cassaigne. A critical analysis of bid pricing models and support tool. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, Piscataway, NJ, 2000.
- [19] D. Pardoe and P. 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.
- [20] R. S. Pindyck and D. L. Rubinfeld. *Econometric Models and Econometric Forecasts; 4th Edition*. Irwin/McGraw-Hill, 1998.
- [21] E. Sodomka, J. Collins, and M. Gini. Efficient statistical methods for evaluating trading agent performance. In *Proc. of the Twenty-Second Nat'l Conf. on Artificial Intelligence*, 2007.
- [22] S. Sun, V. Avasarala, T. Mullen, and J. Yen. PSUTAC: A trading agent designed from heuristics to knowledge. In *Workshop on Trading Agent Design and Analysis at AAMAS*, pages 15–20, New York, July 2004.
- [23] D. Titterton, A. Smith, and U. Makov. *Statistical Analysis of Finite Mixture Distributions*. Wiley, New York, 1985.
- [24] D. Zhang, K. Zhao, C.-M. Liang, G. B. Huq, and T.-H. Huang. Strategic trading agents via market modeling. *SIGecom Exchanges*, 4(3):46–55, 2004.