Tactical and Strategic Sales Management for Intelligent Agents Guided By Economic Regimes

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

ERIM REPORT SERIES *RESEARCH IN MANAGEMENT*	
ERIM Report Series reference number	ERS-2008-061-LIS
Publication	October 2008
Number of pages	23
Persistent paper URL	http://hdl.handle.net/1765/13547
Email address corresponding author	wketter@rsm.nl
Address	Erasmus Research Institute of Management (ERIM) RSM Erasmus University / Erasmus School of Economics Erasmus Universiteit Rotterdam P.O.Box 1738 3000 DR Rotterdam, The Netherlands Phone: + 31 10 408 1182 Fax: + 31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES

RESEARCH IN MANAGEMENT

ABSTRACT AND KEYWORDS	
Abstract	We present a computational approach that autonomous software agents can adopt to make tactical decisions, such as product pricing, and strategic decisions, such as product mix and production planning, to maximize profit in markets with supply and demand uncertainties. Using a combination of machine learning and optimization techniques, the agent is able to characterize economic regimes, which are historical microeconomic conditions reflecting situations such as over-supply and scarcity. We assume an agent is capable of using real-time observable information to identify the current dominant market condition and we show how it can forecast regime changes over a planning horizon. We demonstrate how the agent can then use regime characterization to predict prices, price trends, and the probability of receiving a customer order in a dynamic supply chain environment. We validate our methods by presenting experimental results from a testbed derived from the Trading Agent Competition for Supply Chain Management (TAC SCM). The results show that our agent outperforms traditional short- and long-term predictive methodologies (such as exponential smoothing) significantly, resulting in accurate prediction of customer order probabilities, and competitive market prices. This, in turn, has the potential to produce higher profits. We also demonstrate the versatility of our computational approach by applying the methodology to prediction of stock price trends.
Free Keywords	agent-mediated electronic commerce, dynamic pricing, machine learning, rational decision making, market forecasting
Availability	The ERIM Report Series is distributed through the following platforms: Academic Repository at Erasmus University (DEAR), DEAR ERIM Series Portal Social Science Research Network (SSRN), SSRN ERIM Series Webpage Research Papers in Economics (REPEC), REPEC ERIM Series Webpage
Classifications	The electronic versions of the papers in the ERIM report Series contain bibliographic metadata by the following classification systems: Library of Congress Classification, (LCC) LCC Webpage Journal of Economic Literature, (JEL), JEL Webpage ACM Computing Classification System CCS Webpage Inspec Classification scheme (ICS), ICS Webpage

Tactical and Strategic Sales Management for Intelligent Agents Guided By Economic Regimes

Wolfgang Ketter*, John Collins†, Maria Gini†, Alok Gupta‡, and Paul Schrater†
*Dept. of Decision and Information Sciences, Rotterdam Sch. of Mgmt., Erasmus University
†Dept. of Computer Science and Engineering, University of Minnesota
‡Dept. of Information and Decision Sciences, Carlson Sch. of Mgmt., University of Minnesota
wketter@rsm.nl, {jcollins,gini,schrater}@cs.umn.edu, alok@umn.edu

Abstract

We present a computational approach that autonomous software agents can adopt to make tactical decisions, such as product pricing, and strategic decisions, such as product mix and production planning, to maximize profit in markets with supply and demand uncertainties. Using a combination of machine learning and optimization techniques, the agent is able to characterize *economic regimes*, which are historical microeconomic conditions reflecting situations such as over-supply and scarcity. We assume an agent is capable of using real-time observable information to identify the current dominant market condition and we show how it can forecast regime changes over a planning horizon. We demonstrate how the agent can then use regime characterization to predict prices, price trends, and the probability of receiving a customer order in a dynamic supply chain environment. We validate our methods by presenting experimental results from a testbed derived from the Trading Agent Competition for Supply Chain Management (TAC SCM). The results show that our agent outperforms traditional short- and long-term predictive methodologies (such as exponential smoothing) significantly, resulting in accurate prediction of customer order probabilities, and competitive market prices. This, in turn, has the potential to produce higher profits. We also demonstrate the versatility of our computational approach by applying the methodology to prediction of stock price trends.

1 Introduction

In recent years automated decision support systems have become increasingly sophisticated. Businesses and system designers are developing intelligence in software systems to create software agents that can make autonomous decisions by acting rationally on behalf of human users in a variety of application areas. Some of the examples include procurement ([40]), scheduling and resource management ([20, 8]), and personal information management ([3]). Several researchers have provided guidance for creating agent based architectures for supply chain management (for instance, [42, 12]).

Software agents have the advantage of being able to analyze many more possibilities in shorter timeframes than their human counterparts, but are often limited in their ability to make strategic decisions. In this paper, we develop a computational methodology for a software agent that observes its economic environment and predicts the future economic state of markets it operates in. We are particularly interested in supply-chain environments that are constrained by capacity and materials availability and where market conditions may be characterized qualitatively, for example, by over-supply or scarcity. The environment we consider is a multi-commodity market with highly variable prices.

We describe a computational approach for a sales management agent that uses its awareness of market conditions, that we call *economic regimes*, to make decisions regarding sales strategy over a planning horizon. An agent competes with several other sales agents (human or computational) to fulfill Requests for Quotes (RFQs) from a population of customers. Besides responding to current RFQs, an agent also has to make long-term strategic decisions regarding how much inventory of finished goods and raw material to hold as well as decide production schedules. The agent proposed in this work predicts whether the future economic environment will be more or less favorable. If the future economic environment is likely to be better than the current one, the agent may decide to hold inventory or quote a higher price, and vice-versa. The agent operates in a highly dynamic and competitive environment; therefore, the agent has to be responsive to the signals it receives from the market.

We present a model that predicts future changes in economic regimes (such as oversupply or scarcity) and show how an agent can use this information to make both tactical pricing responses to customer RFQs, and strategic decisions over a planning horizon to maximize long-term profits. We implement the model and demonstrate the utility of our approach in the context of an autonomous agent that is designed to compete in the Trading Agent Competition for Supply Chain Management (TAC SCM)([7]).

In previous research ([24]) we have shown that economic regimes can be identified from historical data. In this paper, we show how machine learning approaches can be used to recognize economic regimes in real-time and predict future transitions in economic regimes. Further, we show how regime identification and prediction can be used to set sales quotas for current and future sales to maximize profit. While predictions about the economic environment are commonly made at the macroeconomic level ([36]), to our knowledge, such predictions are rarely done for microeconomic environments and represent a novel contribution of this research.

Although we primarily focus on the supply-chain trading environment to demonstrate our approach and techniques, our method is general and has application in any domain where rapid decision making is at premium and depends on the economic environment. Examples include agents for automated trading in financial markets, such as the Penn-Lehman Automated Trading Project ([21]), auction-based contracting environments, such as MAGNET ([9]), and other auctions, such as auctions for IBM PCs ([30]) or PDA's on eBay ([13]).

In §2 we review relevant literature and briefly describe the supply-chain trading environment of TAC SCM. In §3 we describe the information that is available in such environments and how strategic and tactical sales decisions can be made in such scenarios. In §4, we describe how to make predictions about future economic regimes. In §5 we apply our methodology to the TAC SCM environment and present experimental results. In §6 we briefly demonstrate the versatility of our approach by applying our methodology to prediction of stock prices. Finally, we conclude with directions for future research.

2 Background and Literature Review

Predicting prices is an important part of the decision process of agents or human decision makers. While the approaches for prediction of prices vary considerably, it is widely recognized that predictions need to exploit the information available in the market and to take its structure into account ([33]). An interesting recent example of such an approach by [22] explored several dynamic pricing algorithms for information goods, where shopbots look for the best price, and pricebots adapt their prices to attract business. Similarly, [45] developed metrics for price prediction in the TAC Classic game ([44]) which focuses on travel and leisure arrangements by travel agents for their clients.

Short-term price prediction has been a focus of several studies where prices move primarily due to demand-side constraints, such as in the electricity market ([35]). However, as [16] note, demand-side price movements are intrinsically linked with supply side movements. [31] show that the ability of decision makers to correctly identify the onset of a new regime can mean the difference between success and failure. Furthermore, they found strong evidence that individuals pay inordinate attention to the signal (price in our case), and neglect diagnosticity (regime dynamics) and transition probability, i.e., the aspects of the system that generate the signal. This results in a tendency to over- or under-react to market conditions. Several researchers have identified the existence and cyclic nature of economic regimes in consumer markets. For example, [14] empirically analyze the degree to which used products cannibalize new product sales for books on Amazon.com. In their study they show that product prices go through different regimes over time. Similarly, [38] analyze how in mature economic markets strategic windows of change alternate with long periods of stability.

In this paper we develop computationally efficient methods to identify and predict economic regimes that can be used by autonomous computational agents. Several researchers have developed methods for identifying models of rational decision-making in autonomous agents from their actions. For example, [34] show that an agent's decisions can be viewed as a set of linear constraints on the space of possible utility (reward) functions. However, the simple incentive structures used in [34] are unlikely to scale for price prediction in the complex economic environment that we consider in this paper. [6] describe a method for predicting the future decisions of an agent based on its past decisions; their approach is based on assuming rational decision makers who maximize their expected utility. However, such methods cannot be applied when the behaviors of competing agents are not directly observable.

Our approach is flexible and can be adapted to environments characterized by a variety of types and quality of information available. We chose to test our approach in an environment characterized by dynamic procurement, production and sales, with limited information regarding the behavior of competitors, the Trading Agent Competition for Supply Chain Management (TAC SCM) ([7]). Before summarizing the literature on TAC SCM, we first briefly describe this competitive environment.

TAC SCM is a market simulation where six autonomous agents compete to maximize profits in a computer-assembly scenario. The simulation takes place over 220 virtual days, each lasting fifteen seconds of real time. TAC SCM agents earn money by selling computers they assemble using several parts that they competitively acquire from suppliers. Each agent has a limited manufacturing capacity to allocate across a set of different products. Each agent must pay to store raw materials and finished-product inventory, and must borrow money to build its initial inventory. The agent with the highest bank balance at the end of the simulation wins.

An agent can produce 16 different types of products that are categorized into three different market segments (low, medium, and high quality products). Demand in each market segment varies randomly during the game. Every day each agent receives a set of *requests for quotes* (RFQs) from several potential *customers*. Each customer RFQ specifies the type of product requested, along with quantity, due date, reserve price, and penalty for late delivery. Each agent may choose to bid on some or all of the day's RFQs. Customers accept the lowest bid that is at or below their reserve price, and notify the winning agent. The agent must ship customer orders on time, or pay the penalty for each day an order is late. Since the environment is essentially a competitive oligopolistic market, actions of each agent significantly affect other agents' profits and strategy.

Several researchers have modeled the probability of receiving an order in TAC SCM for a given offer price, either by estimating it by linear interpolation from the minimum and maximum daily prices ([37]), or by estimating the relationship between offer price and order probability with a linear cumulative density function (CDF) ([2]), or by using a reverse CDF and factors such as quantity and due date ([25]). Some researchers ([46]) have applied a game theoretic approach to set offer prices, using a variation of the Cournot game for modeling the product market. Other researchers ([17]) use fuzzy reasoning to set offer prices. Similar techniques have been used outside TAC SCM to predict offer prices such as in first price sealed bid reverse auctions for IBM PCs' ([30]), PDAs' on eBay ([13]), or in predicting ending prices for a multi-unit online ascending auction ([1]). The problem of allocating finite resources across a set of potential products in a way that maximizes some measure of utility is the well-known "product-mix" problem ([18]). [27] demonstrate a method for predicting future customer demand in TAC SCM. Their approach is specific to TAC SCM, since it depends on knowing the formula by which customer demand is computed.

These methods fail to take into account market conditions that are not directly observable. They are essentially regression models, and do not represent qualitative differences in market conditions. Our method, in contrast, is able to detect and forecast a broader range of market conditions. Regression based approaches (including non-parametric variations) assume that the functional form of the relationship between dependent and independent variables has a consistent structure across the range of market conditions. In contrast, our approach models variability in market conditions and does not assume a functional relationship; this allows detection of changes in relationship between prices and sales over time. In the next section we describe the basic computational requirements to make autonomous strategic and tactical sales decisions in dynamic supply chain environments.

3 Strategic and tactical decisions

We wish to maximize the profit an agent can expect to earn over some reasonable period in the future. Our approach is to treat procurement, production, and sales as separate components with their own decision processes, sharing access to a common store of data and models. This is common in many industries where procurement and production are often driven by relatively long-term forecasts, and sales is expected to move the product it has available to sell (and expected to have in the future) at the best possible price. An agent making sales decision in markets that are affected by price fluctuation needs to make two broad decisions: i) whether to sell or hold inventory; and ii) if the decision is to sell at least part of the inventory, what price it should quote. Holding inventory makes sense when the agent is able to assess that the future is likely to sustain higher prices. On the other extreme, if the agent is holding a large inventory and the future economic outlook looks bleak, the agent should sell down inventory to liquidate it. The decision to hold a certain level of inventory for the future is a strategic decision and setting the price for a given time period is a tactical decision. We next discuss how an autonomous agent can make these decisions.

3.1 Strategic decision – resource allocation:

Sales decisions can be informed by experience from the past, and observations in the present. We first focus on the information visible to the agent at present, which in a manufacturing environment includes the following:

- $\mathcal{C}$ is the set of all available component types. Each component c is needed to produce some subset of products $\mathcal{G}_c$.
- $\mathcal{G}$ is the set of all products that can be manufactured and sold. Each product's components are represented by the set $\mathcal{C}_g$.
- For a day d within the agent's planning horizon h, customer demand is represented by a set $\mathcal{Q}_d$ of customer RFQs. Each $q \in \mathcal{Q}_d$ specifies a product type g_q, a lead time of i_q days, a volume v_q, and a reserve price ρ_q.
- For a day d within the planning horizon h, the agent expects to have 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 $g \in \mathcal{G}$.
- On any given day d, there is an unsold inventory $I'_{d,g}$ 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.

On any given day d, the total demand $D_{d,g}$ for a given good g among Q_d is the total of the requested quantities among requests for good g, $D_{d,g} = \sum_{q \in \mathcal{Q}_d} v_q$.

The *effective demand* $D^{\mathit{eff}}_{d,g}$ is the portion of total demand with reserve prices $\rho_g \geq \mathit{price}_{d,g}$:

$$D^{\mathit{eff}}_{d,g} = \frac{\rho^{\max}_g - \mathit{price}_{d,g}}{\rho^{\max}_g - \rho^{\min}_g} D_{d,g} \tag{1}$$

where $\rho^{\min}_g$ and $\rho^{\max}_g$ are the minimum and maximum reserve prices for good g. For analytical tractability, we assume a uniform distribution of reserve prices ρ_q among customer RFQs $\mathcal{Q}$ for a given good g.[1]

[1] This is the case in TAC SCM

The goal of the agent is to choose a sales quantity $A_{d,g}$ for each product g over each day of the planning horizon h to maximize its expected profit $\Phi = \sum_{d=0}^{h} \sum_{g \in \mathcal{G}} \Phi_{d,g} A_{d,g}$, where $\Phi_{d,g}$ is the discounted profit for day d and $A_{d,g}$ is the quantity of product the agent wishes to sell for good g on day d. The discounted profit is computed as:

$$\Phi_{d,g} = \gamma_d (price_{d,g} - cost(\mathcal{C}_g)) \tag{2}$$

where γ_d is a discount term that can be seen as a rough approximation of inventory holding and opportunity costs. It can also be used to encourage early selling, as a hedge against future uncertainty. The price $price_{d,g}$ for product g on day d will depend on the demand $D_{d,g}$ and the quantity of product $A_{d,g}$ the agent wishes to sell, as well as other factors that we will discuss in §3.2.

The daily production capacity of an agent is F, each unit of good g requires y_g production cycles, and F_m^{commit} is the factory capacity that is committed to manufacture outstanding customer orders that are due on or before a day m days in the future and are not satisfiable by existing finished goods inventory. Now, we can define the agent's objective function of maximizing the total profit, Φ, by choosing appropriate sales quotas $A_{d,g}$:

$$\max \qquad \Phi = \sum_{d=0}^{h} \sum_{g \in \mathcal{G}} \Phi_{d,g} A_{d,g} \tag{3}$$

$$subject\ to: \qquad \forall d, \forall g, A_{d,g} < D_{d,g}^{eff} \tag{4}$$

$$\forall m \in 0..h, \forall c \in \mathcal{C}, \sum_{d=0}^{m} \sum_{g \in \mathcal{G}_c} A_{d,g} \leq I'_{m,c} + \sum_{g \in \mathcal{G}_c} I'_{m,g} \tag{5}$$

$$\forall m \in 0..h, \sum_{g \in \mathcal{G}} y_g \left(\sum_{d=0}^{m} A_{d,g} - I'_{d,g} \right) \leq mF - F_m^{commit} \tag{6}$$

Constraint 4 is the demand constraint. Constraint 5 is the supply constraint over the planning horizon, h, that restricts maximum supply that can be created using the parts in existing inventory. This is conservative, since we are considering goods or their parts to be available at the time we propose to *sell* them, not when we expect to *ship* them. The constraint also ensures that every subset of product types that can share some component is not overcommitted. Constraint 6 is the manufacturing constraint that restricts the sales quantity to what is in the unsold inventory or can be manufactured within the planning horizon.

To appropriately choose sales quotas $A_{d,g}$, we need to set prices. Suppose we can compute the probability of a customer placing an order as a function of price $P(order|price)$. Since the quantity we expect to sell is just the effective demand multiplied by the probability of order at the price we set, we can then express $A_{d,g}$ as:

$$A_{d,g} = P(order|price_{d,g}) D_{d,g}^{eff} = P(order|price_{d,g}) \frac{\rho_g^{\max} - price_{d,g}}{\rho_g^{\max} - \rho_g^{\min}} D_{d,g} \tag{7}$$

Combining (2) with (7), the objective function (3) becomes

$$\max \Phi = \sum_{d=0}^{h} \sum_{g \in \mathcal{G}} \gamma_d \left(price_{d,g} - cost(\mathcal{C}_g) \right) P(order|price_{d,g}) \frac{\rho_g^{\max} - price_{d,g}}{\rho_g^{\max} - \rho_g^{\min}} D_{d,g} \tag{8}$$

Note, even if we assume that the order probability is linear, (8) is at least cubic in $price_{d,g}$. Since (8) is probably unsolvable in real-time, we focus on developing heuristics that can be embedded in automated agents. 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 $price_{d,g}^{clear}$ for the actual sales price into (8). We will explore a way to compute $price_{d,g}^{clear}$ in the next section. However, we first discuss how the strategic sales process guides the tactical decision.

3.2 Tactical decision – sales offer pricing

Once the strategic sales process has determined daily sales quotas, we must set prices that will move those quotas in expectation. This amounts to finding, for each good, the value for $price_{d,g}$ that satisfies (7). We call this $price^{offer}_{d,g}$, and we estimate it by first estimating the market clearing price $price^{clear}_{d,g}$ and using it to locate an order-probability distribution $P(order|price)$ as described in §4.3. The clearing price for the current day is estimated by combining a predicted price with an offset that is computed by observing the market's response to our offers, as follows.

We compute $price^{offer}_{d,g}$ by choosing a target order probability $P^{offer} = \frac{A_{d,g}}{D^{eff}_{d,g}(price^{clear}_{d,g})}$ and finding the corresponding offer price $price^{offer}_{d,g}$ from (7) by solving $P^{offer} = P(order|price^{offer}_{d,g})$. Assuming the market clears once each day, the order volume $O_{d,g}$ is the number of orders placed for good g in response to our offers on the previous day. The market is somewhat unpredictable, so the number of orders we receive may be higher or lower than our expected sales. We compute a refined offer price $price^{offer'}_{d-1,g}$ for the previous day by computing a point $P^{offer'} = \frac{O_{d,g}}{D^{eff}_{d-1,g}(price^{clear}_{d-1,g})}$ on an adjusted probability curve $P^{offer'}(order|price^{offer'}_{d-1,g})$, obtained by translating the curve through P^{offer} to the point $P^{offer'}$. We then use this adjusted curve to compute $price^{offer'}_{d-1,g}$. Figure 1 visualizes this relationship.

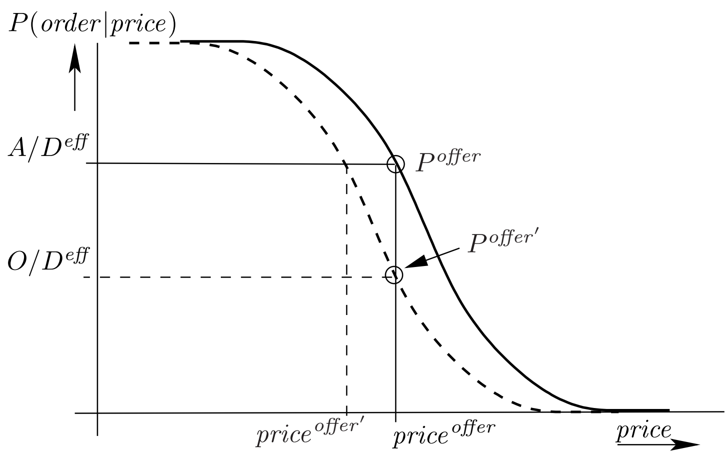

Figure 1: Estimating market price, given order volume O, sales quota A, and an order probability function P.

The difference $diff_{d-1,g} = price^{offer'}_{d-1,g} - price^{offer}_{d-1,g}$ is then used to compute $price^{clear}_{d,g}$ each day, as $price^{clear}_{d,g} = price^{pred}_{d,g} + \delta_{d,g}$ where $price^{pred}_{d,g}$ is the *predicted* market price for product g (see §4.3), and $\delta_{d,g}$ is updated daily as $\delta_{d,g} = \alpha\delta_{d-1,g} + (1-\alpha)diff_{d-1,g}$ for some appropriate value of $\alpha \in [0, 1]$.

Now that we have discussed what needs to be computed by agents, in the next section we discuss a method to detect and predict economic regimes, and the role of economic regimes in the agent's decision making.

4 Supporting Automated Real-time Sales Decisions

Market conditions change over time, affecting an organization's long- and short-term strategies for procurement, production planning, and pricing. Therefore, any automated decision process should account for prevailing and predicted market conditions. Economic theory suggests that market environments exhibit three broadly defined market conditions: scarcity, balanced, and oversupply. A scarcity condition exists when customer demand is more than product supply in the market, a balanced condition when demand is approximately equal to supply, and an oversupply condition when supply far exceeds customer demand. When there is scarcity, prices are higher, so the agent should price more aggressively. In balanced situations, prices are lower and have more spread, so the agent has a range of options for maximizing expected profit. In oversupply situations, prices are lower and the agent should primarily control costs, and therefore either price based on costs, or wait for better market conditions. In previous research ([24]) we have shown that different market conditions can be identified by using economic principles and machine learning techniques. However, in that work, historical data was available, allowing computation of price distributions that were used to compute mean prices and economic regimes. In this research, we focus on whether regimes can be detected and predicted by an agent in real-time when complete historical data and mean prices are not available.

In the remainder of this paper we would like to aggregate price data for different goods. Since prices are likely to have different ranges for different products, we normalize them by dividing the price of a good by the nominal cost of its components and their assembly cost. We define the normalized price for good g on day d as $\mathrm{np}_{d,g} = \frac{price_{d,g}}{nominal\_cost(\mathcal{C}_g)+assembly\_cost_g}$. In the following, for simplicity of notation, we use np instead of $\mathrm{np}_{d,g}$. Note that np can easily convert back to actual price so that strategic price variables can be computed, such as clearing price and offer price, but np allows for easy aggregation of product prices across categories.

We briefly summarize the analysis of historical data to define economic regimes ([24]) as a foundation for the rest of this paper. We start by fitting a Gaussian mixture model (GMM) ([43]) to historical normalized price data. We use the Expectation-Maximization (EM) Algorithm ([11]) to determine the prior probability, $P(\zeta_i)$, of the Gaussian components of the GMM. The density of the normalized price can be written as:

$$p(\mathrm{np}) = \sum_{i=1}^{N} p(\mathrm{np}|\zeta_i)\, P(\zeta_i) \tag{9}$$

where N is the number of Gaussians in the mixture model and $p(\mathrm{np}|\zeta_i)$ is the contribution of the i-th Gaussian to the normalized price density. The real time update version of this calculation is given later in (17).

Using Bayes' rule we determine the posterior probabilities for each Gaussian ζ_i. We then define the posterior probabilities of all Gaussians given the normalized price np as the following N-dimensional vector: $\vec{\eta}(\mathrm{np}) = [P(\zeta_1|\mathrm{np}), P(\zeta_2|\mathrm{np}), \ldots, P(\zeta_N|\mathrm{np})]$. For each observed normalized price np_j we compute the vector of the posterior probabilities, $\vec{\eta}(\mathrm{np}_j)$, which is $\vec{\eta}$ evaluated at each observed normalized price np_j. The intuitive idea of a regime as a recurrent economic condition is captured by discovering price distributions that recur across time periods in the market. We define *regimes* by clustering price distributions over time periods using the k-means algorithm. The clusters found correspond to frequently occurring price distributions with support on contiguous ranges of np. The center of each cluster is a probability vector that corresponds to regime $r = R_k \forall k = 1, \cdots, M$, where M is the number of regimes. Collecting these vectors into a matrix yields the conditional probability matrix $\mathbf{P}(\zeta|r)$.

After we marginalize over all Gaussians ζ_i we obtain the density of the normalized price np dependent on the regime R_k as $p(\mathrm{np}|R_k)$ as:

$$p(\mathrm{np}|R_k) = \sum_{i=1}^{N} p(\mathrm{np}|\zeta_i)\, P(\zeta_i|R_k). \tag{10}$$

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

$$P(R_k|\mathrm{np}) = \frac{p(\mathrm{np}|R_k)\, P(R_k)}{\sum_{k=1}^{M} p(\mathrm{np}|R_k)\, P(R_k)} \quad \forall k = 1, \cdots, M \tag{11}$$

The prior probabilities, $P(R_k)$, of the regimes are determined by a counting process over past data. Mathematical details for computing the optimal number of Gaussians from historical data are presented in the online appendix. We next develop a general computational machinery for real-time regime identification, before demonstrating the effectiveness of our approach in the TAC SCM environment in §5.

4.1 Real time regime identification

Recall that economic regimes characterize prevailing economic conditions. Our goal is to identify the current regime as well as predict future transitions to inform an agent's strategic and tactical decisions. We use the patterns of price changes over time to compute the probabilities of different regimes being present.

In TAC SCM, every day agents are informed of the minimum and maximum prices for each product on the previous day, but they cannot observe sales volume or the distribution of prices. One can use the mid-range normalized price, $\overline{\text{np}}$, the price midway between the minimum and maximum, as a crude approximation for the mean price. However, since observations of minimum and maximum prices are subject to noise, some may be outliers and not representative of the true distribution of the prevailing prices. To motivate our analysis and provide illustrative examples, we use some data from TAC SCM. Figure 2 (left) illustrates an example where mid-range prices do not accurately represent the mean prices. We computed the mean after the game when the entire game data are available. As the example illustrates (especially on day 86, 87, 93, and 110), we observe a spike in the maximum price that biases the mid-range price. To lower the impact of sudden price changes we smooth the minimum and maximum prices using a Brown linear exponential smoother ([5]) with $\beta = 0.5$ to obtain the smoothed minimum $\widetilde{\text{np}}_{d-1}^{min}$ and maximum $\widetilde{\text{np}}_{d-1}^{min}$ normalized prices, from which we compute the smoothed mid-range normalized price $\widetilde{\text{np}}_{d-1}$ as

$$\widetilde{\text{np}}_{d-1} = \frac{\widetilde{\text{np}}_{d-1}^{min} + \widetilde{\text{np}}_{d-1}^{max}}{2} \tag{12}$$

Figure 2: Min, max, mean, mid-range, and smoothed mid-range normalized prices of computers sold every day (left). Realtime identification of daily regime probabilities (right).

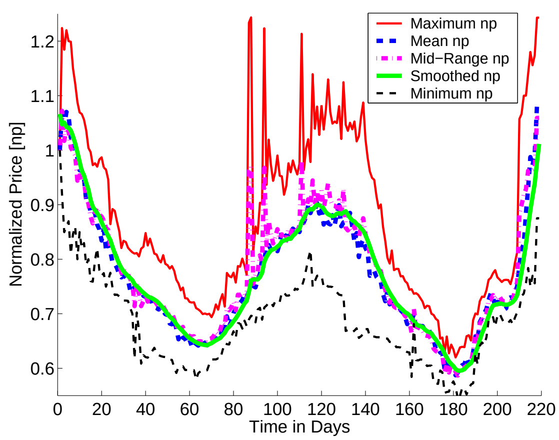

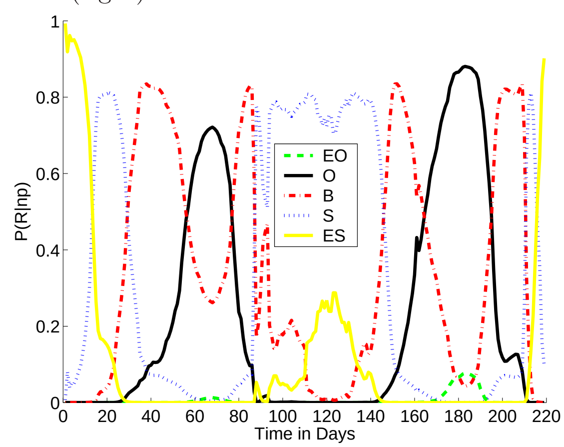

We then select as dominant regime the one which has the highest probability, i.e. $\hat{R}_{max}\, s.t.\, max = \text{argmax}_{1 \leq k \leq M}\, \vec{P}(\hat{R}_k | \widetilde{\text{np}}_{d-1})$. We use $\hat{R}_k$ to denote that R_k is a predicted regime. Figure 2 (left) shows the smoothed mid-range that is used to identify the corresponding regime probabilities in real-time during the game (right). The regimes shown are indicated as EO (Extreme Oversupply), O (Oversupply), B (Balanced), S (Scarcity), and ES (Extreme Scarcity). The graph shows that different regimes are dominant at different time points, and that there are brief intervals during which two regimes are almost equally likely. A correlation analysis of the market parameters to regimes and more details on regime identification and other regime evaluation measures have been reported in [23], [26] and [24].

4.2 Regime prediction

Since the agent's strategic decisions require not just the current regime but also estimates of future regimes, the agent needs a way to predict future regimes. In this section, we describe three different methods for regime prediction. The first is based on exponential smoothing, the second is a Markov prediction process, the last is a Markov correction-prediction process. Each of these methods has different strengths and should be used in different circumstances. The exponential smoother is an ideal candidate to estimate the current regime distribution, since it is more reactive to the current market condition. The Markov prediction process is a good choice for short- and mid-term predictions, while the Markov correction-prediction process is suited for long-term predictions.

4.2.1 Exponential smoother prediction.

First we calculate the minimum smoothed price trend $\widetilde{tr}_{d-1}^{min}$ as

$$\widetilde{tr}_{d-1}^{min} = \frac{\beta}{1-\beta} \cdot (\widetilde{\mathrm{np}}_{d-1}^{min'} - \widetilde{\mathrm{np}}_{d-1}^{min''}) \tag{13}$$

where $\widetilde{\mathrm{np}}'$ and $\widetilde{\mathrm{np}}'\prime$ indicate respectively the singly smoothed and doubly-smoothed normalized prices. $\beta \in (0,1)$ is computational determined parameter between zero and one, that provides computational stability in prediction between the two exponentially smoothed time series. Similarly, we compute the maximum smoothed price trend $\widetilde{tr}_{d-1}^{max}$, and calculate the mid-range trend $\widetilde{tr}_{d-1}$ as the average of $\widetilde{tr}_{d-1}^{min}$ and $\widetilde{tr}_{d-1}^{max}$. Using the trend and yesterday's price estimate (12) we estimate today's and the future daily smoothed prices over the horizon h as:

$$\widetilde{\mathrm{np}}_{d+n} = \widetilde{\mathrm{np}}_{d-1} + (1+n) \cdot \widetilde{tr}_{d-1}, \qquad \forall n = 1, \cdots, h \tag{14}$$

We then obtain the density of $\widetilde{\mathrm{np}}_{d+n}$ dependent on regime $\hat{R}_k$ using (10), and the predicted probability of regime $\hat{R}_k$ dependent on the predicted exponentially smoothed normalized price n days into the future, $\widetilde{\mathrm{np}}_{d+n}$, using (11).

4.2.2 Markov prediction.

We model the short-term prediction of future regimes as a Markov prediction (Markov P) process. The prediction is based on the *last* price measurement. We construct a Markov transition matrix, $\mathbf{T}(r_{d+n}|r_d)$, off-line by a counting process using historical data. 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. We distinguish between two types of Markov predictions: (1) n-day (2) repeated 1-day prediction. The n-day prediction is based on training a separate Markov transition matrix for each day in the planning horizon h, i.e. $\mathbf{T_n}(r_{d+n}|r_d)$, $\forall n = 1, \cdots, h$. The repeated 1-day prediction is done by using multiple times a 1-day prediction matrix, $\mathbf{T_1}(r_{d+1}|r_d)$. We use $\vec{P}(\hat{r}_{d-1}|\{\widetilde{\mathrm{np}}_{d-1}\})$, to indicate a vector of the posterior probabilities of all the regimes on day $d-1$. The prediction of the posterior distribution of regimes n days into the future, $\vec{P}(\hat{r}_{d+n}|\widetilde{\mathrm{np}}_{d-1})$, is done recursively as follows:

1. n-day prediction:

$$\vec{P}(\hat{r}_{d+n}|\widetilde{\mathrm{np}}_{d-1}) = \sum_{r_{d+n}} \cdots \sum_{r_{d-1}} \left\{ \vec{P}(\hat{r}_{d-1}|\widetilde{\mathrm{np}}_{d-1}) \cdot \mathbf{T_n}(r_{d+n}|r_d) \right\} \forall n = 1, \ldots, h \tag{15}$$

2. Repeated 1-day prediction:

$$\vec{P}(\hat{r}_{d+n}|\widetilde{\mathrm{np}}_{d-1}) = \sum_{r_{d+n}} \cdots \sum_{r_{d-1}} \left\{ \vec{P}(\hat{r}_{d-1}|\widetilde{\mathrm{np}}_{d-1}) \cdot \prod_{i=0}^{n} \mathbf{T_1}(r_d|r_{d-1}) \right\}, \forall n = 1, \ldots, h \tag{16}$$

We set the prior regime probability for the first day to 100% extreme scarcity, to represent the starting condition where starting finished product inventories are zero.

4.2.3 Markov correction-prediction.

We model the long-term prediction of future regimes as a Markov correction-prediction (Markov C-P) process, where the prediction part is similar to the Markov prediction described above but takes the entire price history into account. The method is based on two distinct operations:

1. a *correction* (recursive Bayesian update) of the posterior probabilities for the regimes based on the history of measurements of the smoothed mid-range normalized price $\widetilde{\mathrm{np}}$ obtained since the time of the first measurement until the previous day, $d-1$.
2. a *prediction* of regime posterior probabilities for the current day, d. The prediction of the posterior distribution of regimes n days into the future, $\vec{P}(\hat{r}_{d+n}|\{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{d-1}\})$, is done recursively as in the Markov prediction case.

4.3 Prediction of price distribution and trend

Regime prediction is useful to guide an agent's procurement, production, and pricing decisions. Equation 17 shows how to compute the price prediction distribution based on the predicted regime distribution[2] as follows[3]

$$
\begin{aligned}
p(\widehat{\mathrm{np}}_{d+n}|\{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{d-1}\}) &= \sum_{i=1}^{M} p(\mathrm{np}|R_i)\, P(\hat{R}_{i,d+n}|\{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{d-1}\}) \\
&= \sum_{j=1}^{N} \underbrace{\sum_{i=1}^{M} P(\zeta_j|R_i)\, P(\hat{R}_{i,d+n}|\{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{d-1}\})}_{P(\zeta_{j,d+n})}\, p(\mathrm{np}|\zeta_j) \\
&= \sum_{j=1}^{N} P(\zeta_{j,d+n})\, p(\mathrm{np}|\zeta_j), \quad \forall n = 1, \cdots, h \qquad (17)
\end{aligned}
$$

where $P(\hat{R}_{i,d+n}|\{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{d-1}\})$ is an element of the predicted regime probability vector given by (15) or by (16). To obtain a predicted price distribution we sample (17) for every day over the planning horizon h with values for np between 0 and 1.25. Examples of price distributions are shown in Figure 3 (left) and in Figure 4 (left). After sampling the mixture distribution over the set of np values, the distribution is renormalized to sum to one (indicated as p_{norm}). This discretizes the continuous distribution, which simplifies all subsequent probability calculations. For instance the mean of the distribution can be computed as:

$$
E[\widehat{\mathrm{np}}_{d+n}] = \sum_{j=1}^{N} p_{norm}\left(\widehat{\mathrm{np}}_{d+n}(j)|\{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{d-1}\}\right) \cdot \mathrm{np}(j), \qquad \forall n = 1, \cdots, h \qquad (18)
$$

where $\mathrm{np}(j)$ is the sum of all npstarting from np=0.00 to np=1.25 in 0.01 increments. To predict price trends we use also the 10%, 50%, and 90% percentile of the predicted price distribution, which are interpolated from the discretized cumulative distribution.

Figure 3 (left) shows the forecast price density, based on a 1-day Markov matrix, for sample data from TAC SCM (game 3717@tac3 from day 115 to day 135). 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 3 (right) shows the real mean price trend (PT) along with forecast price trends based on the different predictors, the expected mean Markov prediction, the 10%, 50% and the 90% Markov density percentiles, and the exponential smoother.

[2] Recall that M represents the number of regimes and N the number of Gaussians used in the GMM.

[3] We describe the price distribution prediction based on a Markov correction-prediction process, but the same equation holds when using Markov prediction or exponential smoothing.

Figure 3: Predicted price density (left) and predicted price trend (right) using the 1-day Markov matrix for game 3717@tac3 from

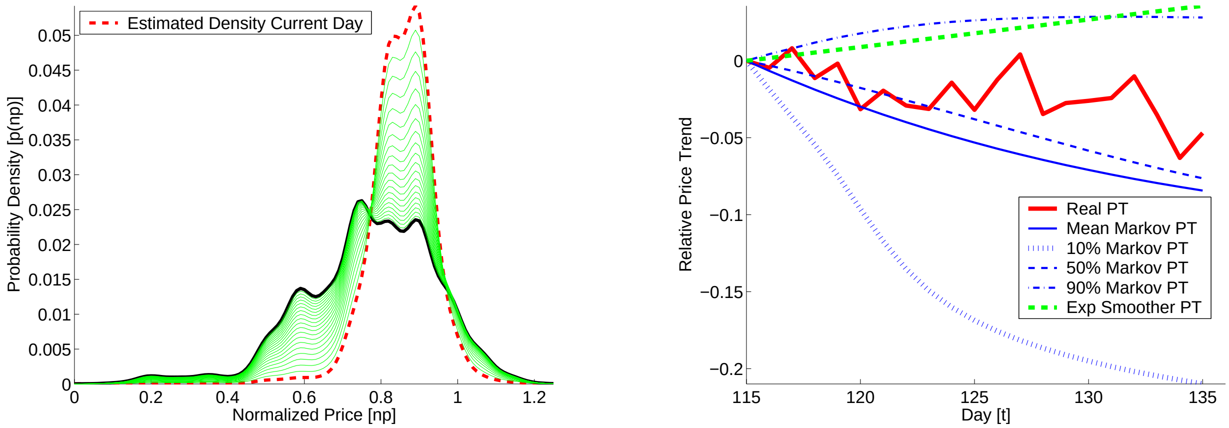

The exponential smoother predictor in this example does not fare well[4], since the smoother is myopic and puts too much weight on recently observed prices. In this case, prior to the prediction day the prices were increasing. The exponential smoother predictor takes this recent slope and extrapolates it into the future, while our Markov prediction method does much better.

Figure 4: Predicted price density (left) and predicted price trend (right) using the repeated n-day Markov prediction.

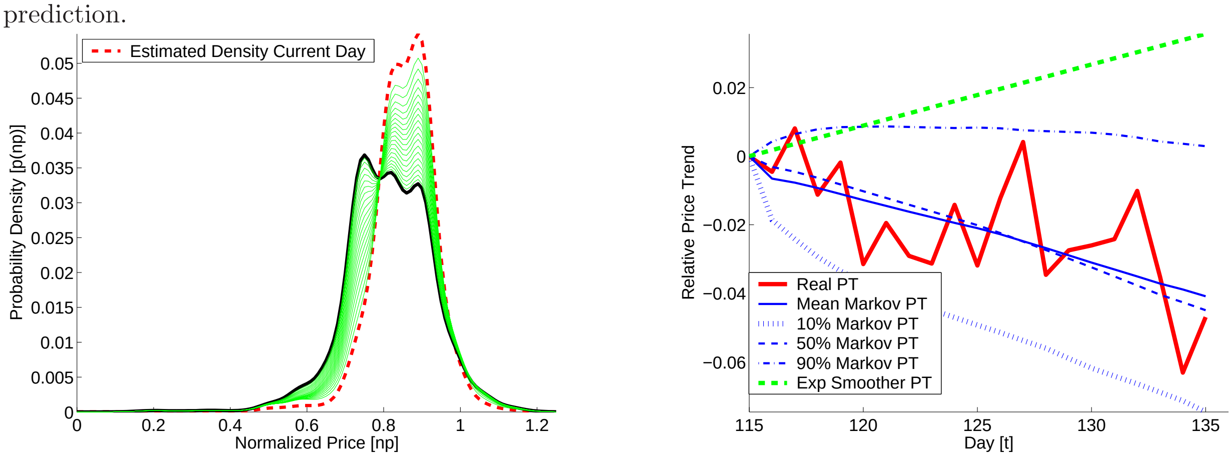

Figure 4 (left) shows the forecast price density based on a n-day Markov prediction for the same game presented above. We observe that the predicted price density shows significantly less variance as compared to using the 1-day Markov matrix. Figure 4 (right) shows the relative price trend. The increased certainty in prediction is reflected in the extremes of the predicted density, since the 10% and 90% percentiles price predictions form a much tighter prediction envelope. Note that the downward shift in actual prices, Figure 4 (right), is captured by the shift of the predicted future price distribution towards the lower prices in the left panel of Figure 4.

[4]Usually the exponential smoother predicts much better (§5), but we use this example to explain one of the advantages of our method.

4.4 Prediction of order probability

The cumulative density function $CDF(\mathrm{np})$ describes the proportion of orders that will be placed at or below the normalized price np. Therefore the probability of a customer order, $P(order|\mathrm{np})$, can be written as: $P(order|\mathrm{np}) = 1 - CDF(\mathrm{np}) = 1 - \int_0^{\mathrm{np}} p(\mathrm{np}')\,\mathrm{dnp}'$

5 Pricing and Sales Decisions in TAC SCM

In this section we evaluate our prediction methods by applying our approach to compute price distribution, price trends, and probability of order in the TAC SCM environment. We begin by presenting compelling results based on historical data, followed by results from real-time experiments.

5.1 TAC SCM - Historical data

In previous research we have shown that economic regimes capture market information that is not directly observable ([24]). There are many hidden attributes in a competitive market, such as the inventory positions and procurement arrangements of the competitors. Our method uses observable historical and current data to guide tactical and strategic decisions. We now demonstrate the practical value of our method by using it with historical data.

5.1.1 Experimental setup.

For our experiments, we used data from 28 games (18 for training and 10 for testing - see online Appendix) played during the semi-finals and finals of TAC SCM 2005. The mix of agents changed during the games, with a total of 12 agents in the semi-finals and 6 in the finals. Since supply and demand vary in each market segment (low, medium, and high) independently of the other segments, our method is applied independently in each market segment.

5.1.2 Price distribution.

We forecast the price density for the next n days into the future, where $p(\widehat{\mathrm{np}}_d)$[5] is the predicted price density for the current day, and $p(\widehat{\mathrm{np}}_{d+n})$ is the predicted price density on the n-th day into the future. In our experiments we chose an horizon $h = 40$. We calculated the expected mean price using (18), and tracked different areas (10%, 50%, and 90%) of the price density curve. We also calculated the expected mean price using the exponential smoother as an input to the regime model to predict the whole price distribution, and the simple exponential smoother to predict prices. We calculated the root mean square error, $RMSE(\widehat{\mathrm{np}}_n, \mathrm{np}_n)$, between the predicted normalized prices $\widehat{\mathrm{np}}_n$ and the actual normalized price, np_n, over a prediction interval of n days in the planning horizon h, averaged across days and games, to determine the accuracy of the price prediction as:

$$RMSE(\widehat{\mathrm{np}}_n, \mathrm{np}_n) = \sqrt{\frac{\sum_{i=1}^{N_G} \sum_{d=1}^{N_D - n} \left(\widehat{\mathrm{np}}_d^{n,i} - \mathrm{np}_d^{n,i}\right)^2}{N_G \cdot (N_D - n)}}, \qquad \forall n = 1, \cdots, h \tag{19}$$

where N_D is the number of days in a TAC SCM game and N_G is the number of test games.

Figure 5 shows the RMS error of the Markov predictors using a repeated 1-day matrix (left) versus the n-day matrix (right) and compares it to the RMS error of the price generated by exponential smoother regime lookup and to the simple exponential smoother. An RMS error of 0.05 corresponds to an average prediction error of 4% and an RMS error of 0.25 corresponds to an average prediction error of 20%. It is clear that as compared to a repeated 1-day, the n-day Markov matrix improves the overall price prediction. Typical approaches for price forecasting utilize exponential smoothing or linear regression methodologies ([37]).

[5] For simplicity of notation we leave out the dependence on historical normalized prices.

The results from our experiments show that while the exponential smoother performs reasonably well for short-term prediction, it is too myopic and even a simple modification where exponential smoothing utilizes regime information as described in §4.2.1 performs significantly better. Further, for long term predictions the Markov price predictors (as described in §4.2) perform significantly better than the exponential smoother using regime information.

Figure 5: RMS price prediction error based on a 1-day (left) vs n-day period (right) Markov matrix.

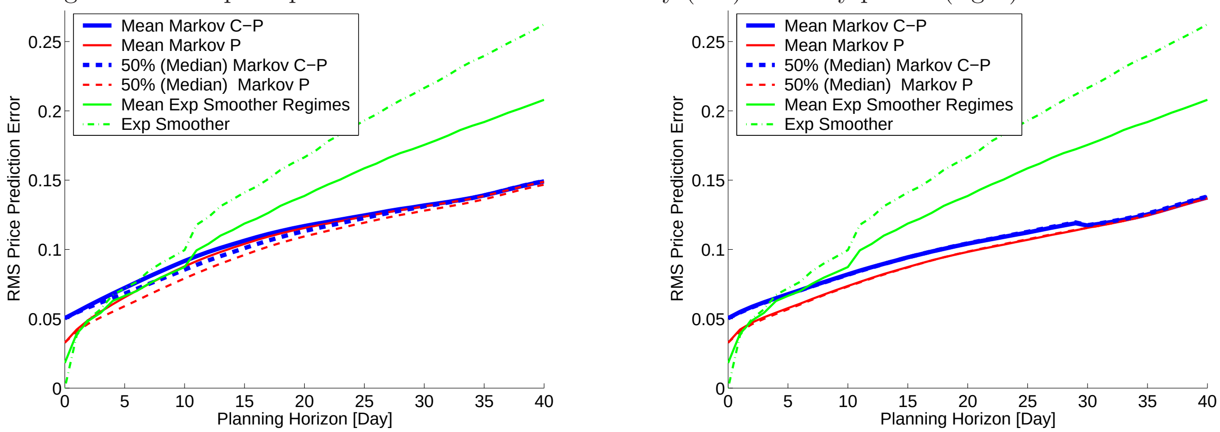

The prices produced by the Markov correction prediction (Markov C-P) are statistically similar to the observed prices since pairwise T-tests testing the equality of the Markov correction-prediction $\widehat{\mathrm{np}}_n$ and actual observed prices np_n failed to reject the null hypothesis at $p = 0.05$.

5.1.3 Price trends.

Besides daily prices, it is important to assess our ability to predict price trends, since they can play a crucial role in sales planning. We computed the estimated price trend $\widehat{Tr}_n$ for every day n over the planning horizon h as follows:

$$\widehat{Tr}_n = sgn(\widehat{\mathrm{np}}_{d+n} - \widehat{\mathrm{np}}_d), \qquad \forall n = 1, \cdots, h \tag{20}$$

where $\widehat{\mathrm{np}}_d$ is the predicted price for the current day, and $\widehat{\mathrm{np}}_{d+n}$ is the predicted price n-days into the future. Recall that the agent has access only to the minimum and maximum prices of the previous day, so it needs a one day forecast to estimate the price on the current day. If $\widehat{Tr}_n$ is positive, then the predicted prices are increasing, while they are decreasing if $\widehat{Tr}_n$ is negative. Figure 6 displays the success rate of price trend sign prediction using a repeated 1-day Markov matrix (left) and a n-day Markov matrix (right). Since the price trend is used for strategic decision making, we calculated the success rate without taking the first five days into account. As the figure demonstrates, Markov correction-prediction predicted the correct trend about 70% of time and dominates the exponential smoothing approach. Also, in general, the n-day Markov predictions performed better than using the 1-day Markov matrix repeatedly.

5.1.4 Order Probability.

Since we estimate the price trends for accepted offers, a direct inverse relationship with order probability can be established. For example, on the normalized price curve a price representing CDF of 10% corresponds to 90% order probability since there is 90% probability that a price higher than that price will be accepted. To test our assertion, we determine, using historical data, how many offers we would have won on each day if we had bid using estimated prices[6]. For our experiment we estimated 2200 (10 games times 220 days) order probability curves for a sample market. Figure 7 shows the results of the experiments for the different

[6] In TAC SCM customers always buy the lowest cost products.

Figure 6: Success-rate of price trend prediction based on 1-day (left) vs. n-day (right) Markov matrix.

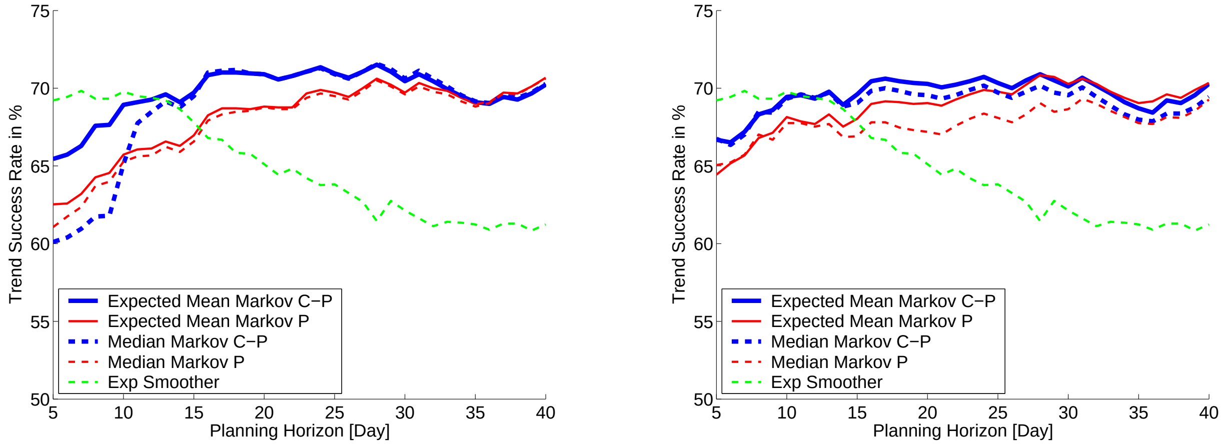

predictors. The y-axis shows the estimated order probability, and the bar graphs show the actual mean order probability and standard deviation. All three predictors estimate the daily order probability well. The ability to estimate the order probability supports the strength of our approach in which we estimate the whole price distribution instead of just mean prices as regression based approaches do.

Figure 7: Daily order probability estimation (mean/std) for the 10th, 25th, 50th, 75th, and 90th percentile using different predictors.

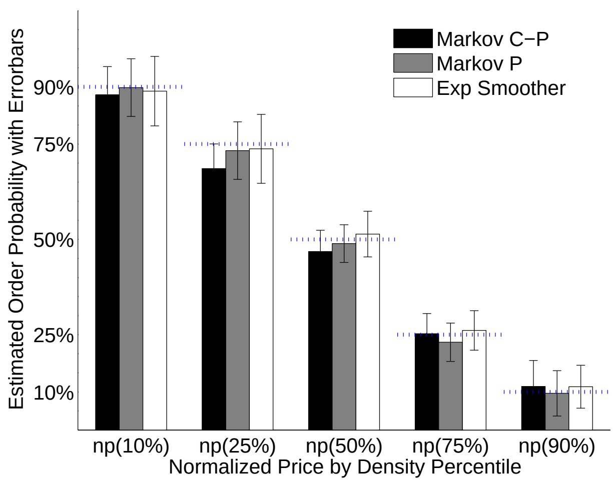

5.2 TAC SCM - Real time

The analysis presented so far demonstrates that our computational approach performs well with historical data. However, to make decisions in real time, the methods have to be dynamic and self-adjusting. We next present the performance of our approach when used by an agent which plays against five other agents in real-time.

5.2.1 Experimental setup.

We implemented all three regime identification and prediction methods, i.e. Markov prediction (Markov P), Markov correction-prediction (Markov C-P), and exponential smoother (ExpS) and tested them in real time in our MinneTAC ([10]) agent. The agents we used for our experiments have been obtained from the TAC SCM agent repository[7]. We selected five of the finalists from the 2006 competition and an agent from the 2005 competition. The agents are: (i) DeepMaize, from University of Michigan; (ii) Maxon, from Xonar Inc; (iii) MinneTAC, from the University of Minnesota; (iv) PhantAgent, from the Politechnica University of Bucharest; (v) RationalSCM, from the Australian National University; and (vi) TacTex, from the University of Texas.

Agent performance in TAC SCM is affected not only by the set of agents playing together but also by random variations in supply, demand, and other market parameters. To compare different variations of our own agent without having to run a very large number of games, we used the controlled server ([41]), which allows for repeatable pseudo-random sequences of any individual market factor or combination of factors. The use of the controlled server removes the profit variability due to agents facing different market conditions, and allowed for testing multiple variations of our agent, one for each set of games. We ran $N_G = 23$ games, each with a different pseudo-random sequence, using the traditional version of MinneTAC, and then ran N_G games with the *same market factors* using different versions of MinneTAC with the three different prediction models described earlier for tactical (order probability calculation when responding to RFQs) and strategic decisions (price and price trend prediction for sales quota and inventory holding decisions). At the strategic level we used two different price prediction methods, the first based on price-following, the second on economic regimes. At the tactical level we used two methods to calculate order probability, one based on a linear interpolation between the estimated minimum and maximum daily prices, the other based on economic regimes.

5.2.2 Real-time results.

Our tests included five sets of twenty-three games each, one set for each different configuration of our MinneTAC agent, using the same twenty-three pseudo-random sequences for each set. For each method we compared the difference in profit and computed the standard error associated with each mean difference. As the primary measure of agent performance we show in Table 1 the mean total profit per agent over a game. Table 1 shows that MinneTAC when playing with this set of agents always comes in fifth; however, the performance of the agent depends upon its ability to acquire raw-material, manufacturing and sales. In this paper, we focus on the sales and consequently we did not change the raw material acquisition and production planning to isolate the effect of regime computations on profits simply by manipulating sales strategies. Therefore, we are only interested in the relative performance of MinneTAC under different sales strategies. The results of the different experiments are as follows:

1. In the first experiment MinneTAC used a linear interpolation to determine the probability of order and an exponential smoother to predict price trends. The a final mean profit is 1.347 million.

2. In the second experiment MinneTAC used again a linear interpolation to determine the probability of order, and economic regimes (based on a repeated 1-day Markov prediction) to predict price trends. The final mean profit was 1.813 million.

3. Experiment three used an exponential smoother to predict prices and a table lookup of previously matched regime probabilities to determine the order probability, median prices and price trends. It had a final mean profit of 1.545 million.

4. Our fourth experiment used regimes for tactical decisions (determination of order probability based on exponentially smoothed predicted regimes) and for strategic decisions (price and price trend prediction based on a repeated 1-day Markov matrix). The final mean profit for this experiment was 2.117 million, the best among the tested configurations.

[7] `http://www.sics.se/tac/showagents.php`

5. Experiment five used a Markov n-day prediction to determine price trends. Its final mean profit was 1.567 million. We expected that online the Markov n-day prediction would outperform the repeated 1-day Markov prediction as reported in §5.1.2, but the outcome of our experiments shows the opposite. The reason could be that off-line we used a separately trained Markov matrix for every day in the planning horizon, but since we have limited time in real-time we used only a 1, 10, and 20 day Markov prediction matrix. Then we performed regime and price density predictions for these three matrices and interpolated the missing prices between them. This assumes that the intermediate prices are linearly related to each other. This is most likely not the case, since we actually expect prices to flatten out further into the future.

Table 1: Experimental results.

	Mean Profit / Standard Deviation (in million)				
Experiment #	1	2	3	4	5
Strategic:	ExpS	Regimes (MP 1-Day)	Regimes (ExpS)	Regimes (MP 1-day)	Regimes (MP n-day)
Tactical:	Linear	Linear	Regimes (ExpS)	Regimes (ExpS)	Regimes (ExpS)
Agent:					
TacTex06	8.752/5.682	8.873/5.600	9.302/5.343	9.205/5.385	9.061/5.331
DeepMaize06F	8.839/4.629	8.713/4.846	8.921/4.733	8.318/4.181	8.652/4.865
PhantAgent06	8.049/5.422	7.991/5.384	8.029/5.425	8.173/5.437	7.953/5.247
Maxon06F	4.243/4.516	3.767/4.288	4.214/4.628	4.019/4.181	3.945/4.396
MinneTAC	1.347/3.703	1.813/4.017	1.545/3.898	2.117/3.764	1.567/3.796
Rational05	0.739/4.912	0.669/4.692	1.032/4.898	1.305/4.527	1.115/4.682

We conducted the Wilcoxon signed rank test ([15, 19]) to assess statistical significance among the first four experiments. The results clearly show that mean profits increase when regimes are used for the purpose of pricing decisions. We conducted Wilcoxon signed rank test ([15, 19]) to assess the statistical significance since the data do not follow a normal distribution due to the state of the games being wildly influenced by random number seeds resulting in many games producing no positive profits by any agent. Note that, since the power of non-parametric tests are smaller than parametric tests, p-values smaller than 0.10 are considered adequate for statistical significance. The result of the tests conclude that while there is no statistical difference in profits between experiment 3 and experiment 5 as compared to experiment 1; the profits are significantly higher in experiment 2 ($p = 0.0523$) and experiment 4 ($p=0.0061$) as compared to experiment 1. We further tested the difference in profits between experiment 4 and experiment 2, to see whether using regimes at the tactical level (i.e., using regimes to predict daily sales prices and order probability to optimize sales strategy) is beneficial as compared to using linear interpolation. The results indicated that the profits are significantly higher in experiment 4 ($p=0.0593$) as compared to experiment 2.

6 Financial markets

In the following we demonstrate the versatility of our approach by applying it to the prediction of price trends in the stock market. An investor could use this approach to decide whether to keep a stock, buy more, or sell in time to make a profit. In terms of the stock market the regimes may reflect "bear" market or "bull" market. As an example we present here the analysis of the stock of General Electric. Note that our goal here is not to compare our approach with other stock prediction approaches, but simply to demonstrate the proof of concept on data outside the TAC SCM domain.

Stock market prices are characterized by a time series, and when we perform the regime training we need to pick a continuous price stream, as opposed to TAC SCM where we randomly pick training games of a pool of games. We obtained the stock market data from the Yahoo finance[8] service. The left side of Figure 8 displays the time series of our training price data from October 1st, 2005 until December 31st, 2005 and the right side shows the price distribution estimation using our approach with GMM ([24]).

[8]Yahoo finance: http://finance.yahoo.com/

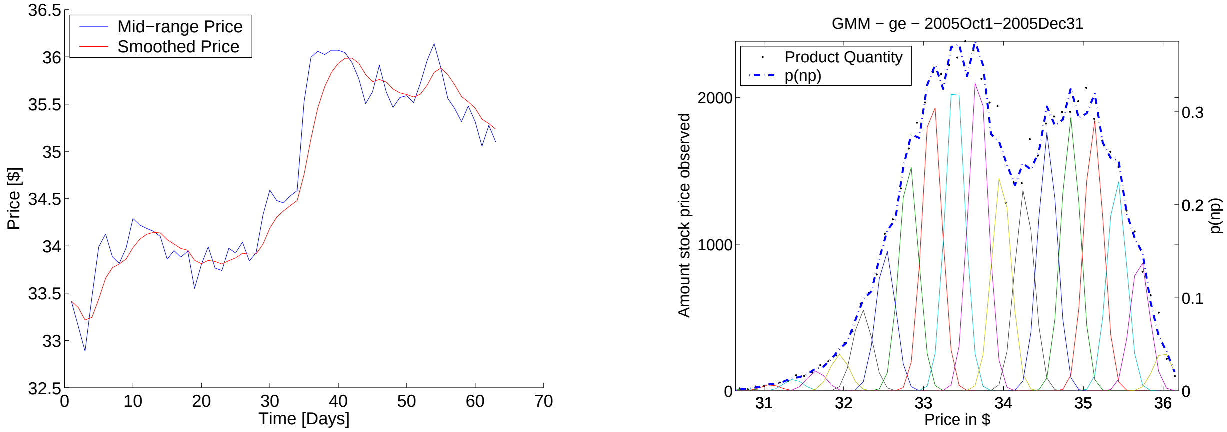

Figure 8: Historical prices from Oct 1st 2005 until Dec 31st 2005 (left) and the appropriate GMM (right).

Figure 9 (left) shows the learned regime probabilities over price. We experimented with different number of regimes on different stocks and found that 5 regimes results in the highest success-rate of price trend predictions. Figure 9 (right) displays the time series of our testing set. We predicted and captured the actual GE stock prices from January 1st 2006 until September 26th 2006 to demonstrate predictive validity of our approach.

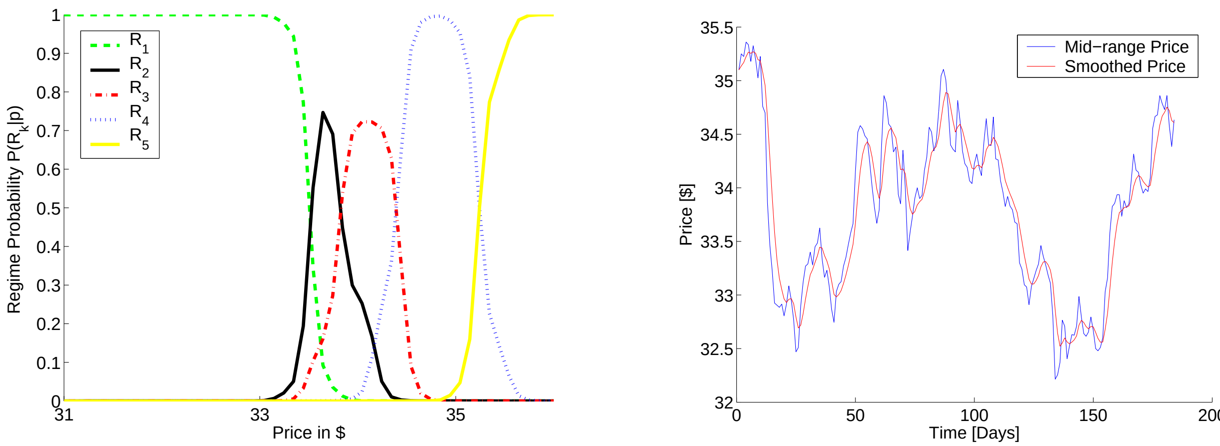

Figure 9: Learned regime distributions (left) and test data from Jan 1st 2006 until Sep 26th 2006 (right).

While an extensive analysis of these results is out of the scope of this paper, Figure 10 presents the success-rate of price trend predictions using a 1-day (left) and a n-day (right) Markov transition matrix. We observe that in the prediction set the accuracy of price trend prediction is above 65% for all planning horizons except for exponential-smoothing based approach and above 75% for longer planning horizons.

7 Conclusions and Future Work

We proposed a versatile computational method based on both historical and observable data that can be used for tactical and strategic economic decision making by automated agents. The approach is based on fundamental economic intuition, recognizing prevailing and predicted economic environments, or regimes, for making pricing and sales decisions. The computational process is completely data driven and no explicit

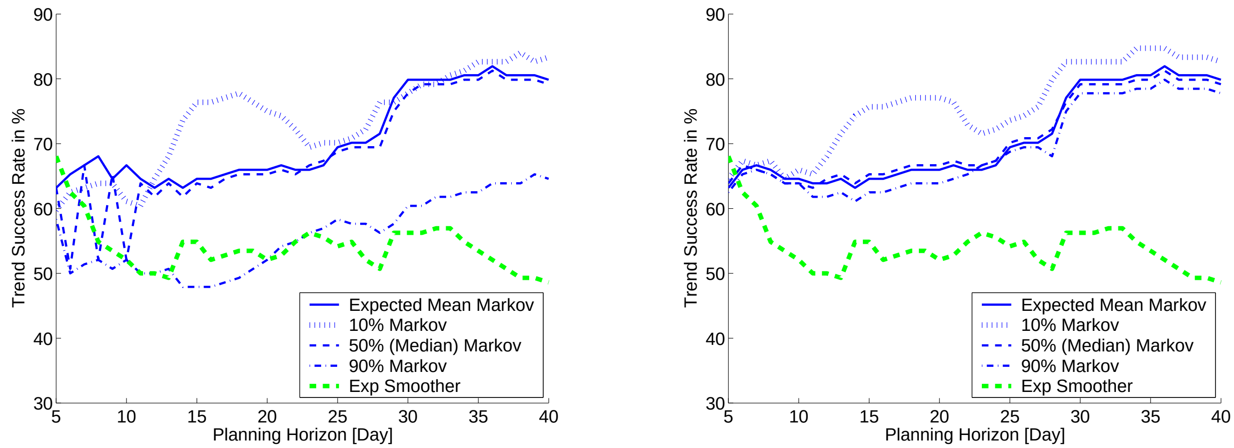

Figure 10: Success-rate of price trend prediction based on 1-day (left) vs. n-day (right) period Markov matrix.

classification of the market structure (monopoly vs competitive, etc.) is needed. A regime encapsulates a complete set of market parameters, with their appropriate range tailored to a specific market condition, thereby reducing the dimensionality of the parameter space. This results in a fast computational approach. Economic regimes provide comparatively more degrees of freedom than ordinary regression based approaches, since the full price distribution is available for decision making. Availability of complete distributions and their trends allows a decision maker to choose an appropriate level of risk, and supports estimation of other useful metrics such as order probabilities. Economic regimes are especially suited to make predictions in non stationary environments where supply-demand relationship is highly dynamic. Economic regimes also provide an opportunity for niche learning, i.e., an agent is able to apply different approaches and actions when specific regimes are dominant. We presented three different algorithms for dynamic identification of regimes and for prediction of regime distribution over a planning horizon. We presented principles and algorithms that use knowledge of current and future regime distributions to facilitate tactical decision making, such as calculation of customer offer prices, and strategic decision making, such as allocation of resources over a planning horizon.

Our approach uses the complete price distribution instead of point-estimates of prices to account for the impact of price variance on decision making. This allows an agent to avoid over-committing to risky decisions. In future, we intend to apply our method in other domains where predicting price distributions maybe fruitful, including domains such as Amazon.com, eBay.com, energy markets, and in financial applications. We also plan to apply machine learning methods to map economic regimes to internal operational regimes and operational regimes to actions, such as procurement and production scheduling

Appendix: Algorithms

Historical data

For our experiments, we used historical data from a set of 28 games (18 for training[9] and 10 for testing[10]) played during the semi-finals and finals of TAC SCM 2005.

[9] 3694@tac3, 3700@tac3, 4229@tac4, 4234@tac4, 7815@tac5, 7821@tac5, 5638@tac6, 5639@tac6, 3719@tac3, 3720@tac3, 3721@tac3, 3722@tac3, 3723@tac3, 4255@tac4,4256@tac4, 4257@tac4, 4258@tac4, 4259@tac4 – To obtain the complete path name append .sics.se to each game number.

[10] 3697tac3, 4235tac4, 7820tac5, 5641tac6, 3717tac3, 3718tac3, 3724tac3, 4253tac4, 4254tac4, 4260tac4

Table 2: Summary of the mathematical notation used in the paper.

Symbol	Definition
$\mathcal{C}$	Set of all available component types
$\mathcal{G}$	Set of all goods (product types)
d	Current day
$D_{d,g}$	Customer demand for good g on day d
$D_{d,g}^{eff}$	Effective customer demand for good g on day d
Φ	Total profit
$A_{d,g}$	Allocated sales quota for good g on day d
F	Factory capacity
h	Planning horizon
np	Normalized price
$\overline{\text{np}}$	Mid-range normalized price
$\widetilde{\text{np}}^{min}$	Smoothed minimum normalized price
$\widetilde{\text{np}}^{max}$	Smoothed maximum normalized price
$\widetilde{\text{np}}$	Smoothed mid-range normalized price
α	Smoothing coefficient
$p(\text{np})$	Density of the normalized price
GMM	Gaussian Mixture Model
N	Number of Gaussians of the GMM
$p(\text{np}\|\zeta_i)$	Density of the normalized price, np, given i-th Gaussian of the GMM
$P(\zeta_i)$	Prior probability of i-th Gaussian of the GMM
$P(\zeta_i\|\text{np})$	Posterior probability of the i-th Gaussian of the GMM given a normalized price np
$\vec{\eta}(\text{np})$	N-dimensional vector with posterior probabilities, $P(\zeta_i\|\text{np})$, of the GMM
M	Number of regimes
R_k	k-th regime, $k = 1, \cdots, M$
$\hat{R}_k$	predicted k-th regime, $k = 1, \cdots, M$
$\mathbf{P}(\zeta\|r)$	Conditional probability matrix (N rows and M columns) resulting from k-means clustering
$p(\text{np}\|R_k)$	Density of the normalized price np given regime R_k
$P(R_k\|\text{np})$	Probability of regime R_k given normalized price np
$P(order\|\text{np})$	Probability of order given a normalized price np
$\mathbf{T}$	Markov transition matrix

Determination of the optimal number of Gaussians for the GMM

We developed an algorithm (see Figure 11) to find the optimal number of Gaussians in the GMM. The algorithm iterates from 1 to N Gaussian components and for each set of Gaussians it fits a GMM to all the historical normalized price data from the training set. New normalized price samples are generated from each fitted GMM model via Monte-Carlo sampling, with the number of new samples matching the original data size. Price histograms are generated using the same bins for the original and sampled data, and are compared with the help of the KL-divergence ([29, 28])[11]. For each set of Gaussians we iterate the resampling and the computation of the KL-divergence. Finally we calculate the mean KL-divergence of all

[11] With the KL-divergence we are able to measure the closeness of two distributions. If the two distributions are exactly the same, then the KL-divergence is zero. A more detailed discussion of the KL-divergence can be found in [24].

the sets of Gaussians. The set with the minimum mean KL-divergence is the set that most closely reproduces the original distribution and is optimal in that sense.

```
Inputs:
    pnp_avg: original normalized price density
    maxNumGauss: the maximum number of Gaussians
    maxFits: iterations of GMM fitting
    NP: set of all normalized prices used for training
    numNP: length of NP
Output:
    optNumGauss: the optimal number of Gaussians
Process variables:
    GMM: Gaussian mixture model
    pnp_samp: sampled estimated normalized price density
    KL: KL divergence
    KL_avg: average KL divergence
Process:
    for comp = 1 until maxNumGauss
        for fits = 1 until maxFits
        GMM = Expectation_Maximization(NP, comp)
        pnp_samp = Monte_Carlo_Sampling(GMM, numNP)
        KL(comp, fits) = KL_divergence(pnp_avg, pnp_samp)
        end
    KL_avg(comp) = mean(KL(comp))
    end
    Index_KL_min = min(KL_avg)
    optNumGauss = KL_avg(Index_KL_min)
    return optNumGauss
```

Figure 11: Algorithm to find the optimal number of Gaussians in a GMM.

The results of the optimization algorithm are in Figure 12 (left), where the mean KL-divergence of 10 fits for 4 to 25 Gaussians and the corresponding standard deviations are plotted. The KL-divergence values for 1 to 3 Gaussians are not displayed since they are too large to fit. The mean KL-divergence for one Gaussian is 2.64, for two is 0.58, and for three is 0.44.

The price density function, $p(\text{np})$, estimated by the GMM with 16 components for a sample market is shown in Figure 12 (right). Even though the optimal number of Gaussians for this sample market is 24, we can see that the GMM with 16 Gaussians fits well the data. For $N = 16$ Gaussians the KL-divergence value is around 0.01, which is small enough to have a good fit to the data.

The number of Gaussians should reflect a balance between accuracy and computational overhead. By accuracy we mean predicted accuracy, which is not the same as fit accuracy. Creating a model with a very good fit to the observed data does not necessarily translate into good predictions. If the model has too many degrees of freedom the likelihood of overfitting the data is great ([32], [39]). [47] and [4] used a similar approach to select an appropriate model with the help of KL-divergences.

)

Figure 12: Mean KL-divergence (left) and price density function (right). The mean KL-divergence is shown for 10 fits of 4 to 25 Gaussians. The standard deviation is too small to be visible. The price density function, $p(np)$, (right) is estimated using 16 Gaussian components. The left y-axis represents the quantity of goods. Data are from 18 games from the semi-finals and finals of TAC SCM 2005.

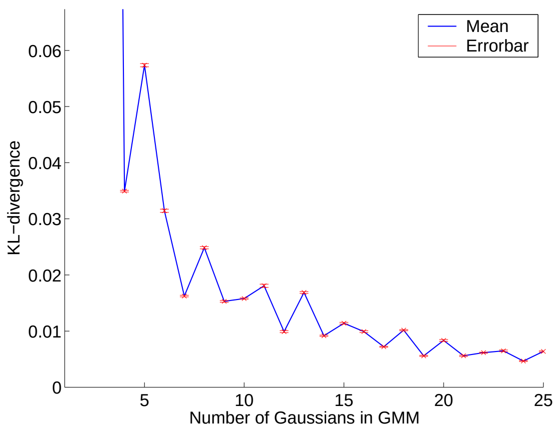

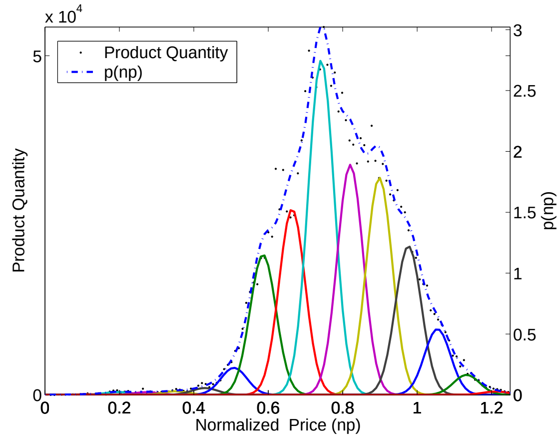

References

[1] R. Bapna, P. Goes, A. Gupta, and G. Karuga. Optimal Design of the Online Auction Channel: Analytical, Empirical and Computational Insights. *JOC*, 20(3):345–355, 2008.

[2] 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. Comp. Logic*, 4(3):29–37, 2004.

[3] P. Berry, K. Conley, M. Gervasio, B. Peintner, T. Uribe, and N. Yorke. Deploying a personalized time management agent. In *Proc. of the Fifth Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, Hakodate, Japan, May 2006.

[4] A. Beygelzimer and I. Rish. Approximability of probability distributions. In *Proc. of the 17th Conf. on Neural Information Processing Systems*, Vancouver, BC, Canada, December 2003.

[5] R. Brown, R. Meyer, and D. D'Esopo. The fundamental theorem of exponential smoothing. *Operations Research*, 9(5):673–687, 1961.

[6] U. Chajewska, D. Koller, and D. Ormoneit. Learning an agent's utility function by observing behavior. In *Proc. of the 18th Int'l Conf. on Machine Learning*, pages 35–42, Lafayette, June 2001.

[7] 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, November 2005.

[8] J. Collins, C. Bilot, M. Gini, and B. Mobasher. Decision processes in agent-based automated contracting. *IEEE Internet Computing*, 5(2):61–72, March 2001.

[9] J. Collins, W. Ketter, and M. Gini. A multi-agent negotiation testbed for contracting tasks with temporal and precedence constraints. *Int'l Journal of Electronic Commerce*, 7(1):35–57, 2002.

[10] J. Collins, W. Ketter, and M. Gini. Flexible decision control in an autonomous trading agent. *Electronic Commerce Research and Applications*, Forthcoming 2008.

[11] 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.

[12] M.S. Fox, M. Barbuceanu, and R. Teigen. Agent-Oriented Supply-Chain Management. *International Journal of Flexible Manufacturing Systems*, 12(2):165–188, 2000.

[13] 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.

[14] A. Ghose, M. Smith, and R. Telang. Internet exchanges for used books: An empirical analysis of product cannibalization and welfare impact. *Information Systems Research*, 17(1):3–19, 2006.

[15] J. Gibbons. Nonparametric statistical inference. *Technometrics*, 28(3):275, 1986.

[16] J.A. Gray and D.E. Spencer. Price Prediction Errors and Real Activity: A Reassessment. *Economic Inquiry*, 28(4):658–681, 1990.

[17] Minghua He, Alex Rogers, Xudong Luo, and Nicholas Jennings. Designing a successful trading agent for supply chain management. In *Proc. of the Fifth Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1159–1166, 2006.

[18] F. Hillier and G. Lieberman. *Introduction to Operations Research*. McGraw-Hill, 1990.

[19] M. Hollander and D. Wolfe. Nonparametric statistical methods. *Journal of the American Statistical Association*, 95(449):333, 2000.

[20] I2. Next-generation planning. http://i2.com/solution_library/ng_planning.cfm, 2006.

[21] M. Kearns and L. Ortiz. The Penn-Lehman Automated Trading Project. *IEEE Intelligent Systems*, pages 22–31, 2003.

[22] J. Kephart, J. Hanson, and A. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, 2000.

[23] W. Ketter, J. Collins, M. Gini, A. Gupta, and P. Schrater. A Computational Approach to Predicting Economic Regimes in Automated Exchanges. In *Workshop on Information Technologies and Systems*, pages 147–152, Las Vegas, USA, December 2005.

[24] W. Ketter, J. Collins, M. Gini, A. Gupta, and P. Schrater. Detecting and Forecasting Economic Regimes in Multi-Agent Automated Exchanges. *Decision Support Systems*, in publication, 2008.

[25] W. Ketter, E. Kryzhnyaya, S. Damer, C. McMillen, A. Agovic, J. Collins, and M. Gini. MinneTAC sales strategies for supply chain TAC. In *Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1372–1373, New York, July 2004.

[26] Wolfgang Ketter. *Identification and Prediction of Economic Regimes to Guide Decision Making in Multi-Agent Marketplaces*. PhD thesis, University of Minnesota, Twin-Cities, USA, January 2007.

[27] C. Kiekintveld, M. Wellman, S. Singh, J. Estelle, Yevgeniy Vorobeychik, Vishal Soni, and Matthew Rudary. Distributed feedback control for decision making on supply chains. In *Int'l Conf. on Automated Planning and Scheduling*, pages 384–392, Whistler, BC, Canada, June 2004.

[28] S. Kullback. *Information Theory and Statistics*. Dover Publications, New York, 1959.

[29] S. Kullback and R. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22:79–86, 1951.

[30] 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, January 2003.

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

[32] T. Mitchell. *Mchine Learning*. McGraw-Hill, 1997. ISBN: 0071154671.

[33] J.F. Muth. Rational Expectations and the Theory of Price Movements. *Econometrica*, 29(3):315–335, 1961.

[34] A. Ng and S. Russell. Algorithms for inverse reinforcement learning. In *proc. Int'l Conf. on Machine Learning*, pages 663–670, Palo Alto, June 2000.

[35] FJ Nogales, J. Contreras, AJ Conejo, and R. Espinola. Forecasting next-day electricity prices by time series models. *IEEE Trans. on Power Systems*, 17(2):342–348, 2002.

[36] D. Osborn and M. Sensier. The prediction of business cycle phases: financial variables and international linkages. *National Institute Econ. Rev.*, 182(1):96–105, 2002.

[37] 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.

[38] K. Pauwels and D. Hanssens. Windows of Change in Mature Markets. In *European Marketing Academy Conf.*, Braga, Portugal, May 2002.

[39] S. Russell and P. Norvig. *Artificial Intelligence - A Modern Approach; 2nd Edition*. Prentice Hall, 2002. ISBN: 0137903952.

[40] T. Sandholm. Expressive commerce and its application to sourcing. In *Proc. of the Twenty-First Nat. Conf. on Artificial Intelligence*, pages 1736–1743, Boston, USA, July 2006.

[41] E. Sodomka, J. Collins, and M. Gini. Efficient statistical methods for evaluating trading agent performance. In *Proc. of the Twenty-Second Nat. Conf. on Artificial Intelligence*, pages 770–775, 2007.

[42] J.M. Swaminathan, S.F. Smith, and N.M. Sadeh. Modeling supply chain dynamics: A multiagent approach. *Decision Sciences*, 29(3):607–632, 1998.

[43] D. Titterington, A. Smith, and U. Makov. *Statistical Analysis of Finite Mixture Distributions*. Wiley, New York, 1985.

[44] M. Wellman, A. Greenwald, and P. Stone. *Autonomous Bidding Agents*. MIT Press, 2007.

[45] M. Wellman, D. Reeves, K. Lochner, and Y. Vorobeychik. Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research*, 21:19–36, 2004.

[46] D. Zhang, K. Zhao, C. Liang, G. Huq, and T. Huang. Strategic trading agents via market modeling. *SIGecom Exchanges*, 4(3):46–55, 2004.

[47] X. Zhang and W. Cheung. Learning global models based on distributed data abstractions. In *Proc. of the Nineteenth Int'l Joint Conf. on Artificial Intelligence*, pages 1645–1646, Edinburgh, Scottland, August 2005.

Publications in the Report Series Research* in Management

ERIM Research Program: "Business Processes, Logistics and Information Systems"

2008

An Evolutionary Framework for Determining Heterogeneous Strategies in Multi-Agent Marketplaces
Alexander Babanov, Wolfgang Ketter and Maria Gini
ERS-2008-002-LIS
http://hdl.handle.net/1765/10972

Choosing between Auctions and Negotiations in Online B2B Markets for IT Services: The Effect of Prior Relationships and Performance
Uladzimir Radkevitch, Eric van Heck and Otto Koppius
ERS-2008-004-LIS
http://hdl.handle.net/1765/11288

Key Issues in Expansion of End-User Mobile Communication in China
Sunanda Sangwan, Guan Chong and Louis-Francois Pau
ERS-2008-011-LIS
http://hdl.handle.net/1765/11762

Some Comments on the Question Whether Co-Occurrence Data Should Be Normalized
Ludo Waltman and Nees Jan van Eck
ERS-2008-014-LIS
http://hdl.handle.net/1765/11810

Bibliometric Mapping of the Computational Intelligence Field
Nees Jan van Eck and Ludo Waltman
ERS-2008-015-LIS
http://hdl.handle.net/1765/11811

Demand Management Opportunities in E-fulfillment: What Internet Retailers Can Learn from Revenue Management
Niels Agatz, Ann Campbell, Moritz Fleischmann, Jo van Nunen and Martin Savelsbergh
ERS-2008-021-LIS
http://hdl.handle.net/1765/12244

Time Slot Management in Attended Home Delivery
Niels Agatz, Ann Campbell, Moritz Fleischmann and Martin Savelsbergh
ERS-2008-022-LIS
http://hdl.handle.net/1765/12245

Mobile Service Affordability for the Needy, Addiction, and ICT Policy Implications
L-F Pau
ERS-2008-023-LIS
http://hdl.handle.net/1765/12246

CBPRS: A City Based Parking and Routing System
J.L. Boehlé, L.J.M. Rothkrantz and M. van Wezel
ERS-2008-029-LIS
http://hdl.handle.net/1765/12467

Coping with Costly Bid Evaluation in Online Reverse Auctions for IT Services
Uladzimir Radkevitch, Eric van Heck and Otto Koppius
ERS-2008-039-LIS
http://hdl.handle.net/1765/12871

Multi Agent Systems in Logistics: A Literature and State-of-the-art Review
Niels Lang, Hans M. Moonen, F. Jordan Srour and Rob A. Zuidwijk
ERS-2008-043-LIS
http://hdl.handle.net/1765/12902

Generalizing the h- and g-indices
Nees Jan van Eck and Ludo Waltman
ERS-2008-049-LIS
http://hdl.handle.net/1765/13210

Applying Revenue Management to the Reverse Supply Chain
Mark Ferguson, Moritz Fleischmann and Gilvan C. Souza
ERS-2008-052-LIS
http://hdl.handle.net/1765/13211

Ergonomics Contributions to Company Strategies
Jan Dul and W. Patrick Neumann
ERS-2008-058-LIS
http://hdl.handle.net/1765/13213

Can the Office Environment Stimulate a Manager's Creativity?
Canan Ceylan, Jan Dul and Serpil Aytac
ERS-2008-059-LIS
http://hdl.handle.net/1765/13215

Tactical and Strategic Sales Management for Intelligent Agents Guided By Economic Regimes
Wolfgang Ketter, John Collins, Maria Gini, Alok Gupta, and Paul Schrater
ERS-2008-061-LIS
http://hdl.handle.net/1765/13547

---

* A complete overview of the ERIM Report Series Research in Management:
https://ep.eur.nl/handle/1765/1

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship