Product Pricing in TAC SCM using Adaptive Real-Time Probability of Acceptance Estimations based on Economic Regimes

Alexander Hogenboom
Erasmus Sch. of Econ.
Erasmus University Rotterdam
287485ah@student.eur.nl

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

Jan van Dalen
Rotterdam Sch. of Mgmt.
Erasmus University Rotterdam
jdalen@rsm.nl

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

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

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

Abstract

Dynamic product pricing is a vital, yet non-trivial task in complex supply chains – especially in case of limited visibility of the market environment. We propose to differentiate product pricing strategies using economic regimes. In our approach, we use economic regimes (characterizing market conditions) and error terms (accounting for customer feedback) to dynamically model the relation between available data and parameters of double-bounded log-logistic distributions assumed to be underlying daily offer prices. Given the parametric estimations of these price distributions, we then estimate offer acceptance probabilities using a closed-form mathematical expression, which is used to determine the price yielding a desired quota. The approach is implemented in the MinneTAC trading agent and tested against a price-following product pricing method in the TAC SCM game. Performance significantly improves. More customer orders are obtained against higher prices and profits more than double.

Introduction

Supply chains are ubiquitous in today's global economy. Raw materials are converted into products and distributed to final users through these complex logistics systems (Ghiani, Laporte, and Musmanno 2004). Every entity adds value to the final product and fulfills a function within the supply chain. Effective Supply Chain Management (SCM), focussing on more flexible and dynamic relationships between entities in the supply chain, is vital to the competitiveness of manufacturers within this chain. SCM can yield this effect, as it enables these manufacturers to respond to changing market demands in a timely and cost effective manner (Collins et al. 2005). Hence, effective SCM can improve the agility of entities within a supply chain. One of the challenges in this context is dynamic product pricing. When flexible and dynamic relationships between supply chain entities stimulate manufacturers to compete for customer orders, optimal product prices are to be determined while accounting for numerous aspects, such as competitors' strategies or market conditions (Dasgupta and Hashimoto 2004; Li, Giampapa, and Sycara 2004; Saha, Biswas, and Sen 2005; Sohn, Moon, and Seok 2009). This requires a sophisticated dynamic product pricing approach, which may be facilitated by automated decision support systems.

In real-time applications, the problem is that aspects relevant to the product pricing process may not be (fully) observable. This is also the case in the Trading Agent Competition for Supply Chain Management (TAC SCM) (Collins et al. 2005), which has been organized since 2002 in order to promote and encourage high quality research into trading agents in supply chains. The game stimulates research with respect to more flexible and dynamic supply chain practices. In TAC SCM, several manufacturers compete in a component procurement market and in a sales market where assembled PCs are sold through reverse auctions where they can bid on requests for quotes (RFQs). The TAC SCM market environment is only partially observable. As dynamic product pricing is a vital, yet non-trivial task in case of limited visibility of the market environment, we investigate how dynamically differentiating product pricing decisions using estimations of economic regimes (Ketter et al. 2006; 2007; 2009) can contribute to profit maximization. We validate our novel approach in TAC SCM.

Related Work

When pricing products, several aspects need to be accounted for. First of all, an approximation of the probability of acceptance of offers can be used, as done in (Dasgupta and Hashimoto 2004) and (Walsh et al. 2008). Analyzing offer prices and their associated estimated probabilities of acceptance is rather intuitive in a product pricing process, because this can help a seller assess how sales targets can be met. In (Dasgupta and Hashimoto 2004), products are priced using a dynamic pricing algorithm which considers an estimated distribution of the buyer reservation price for products of a seller. Reversing the cumulative form of this distribution yields a function expressing the proportion of buyers willing to pay the seller a specified price, which can also be interpreted as the probability that a customer accepts an offered price. This function can subsequently be used to determine the price expected to yield a specified sales quota. In (Walsh et al. 2008), estimated distributions of buyer's private values are used in a similar way.

Another way of modeling acceptance probabilities is by using linear regression on data points representing recent prices offered, along with the resulting acceptance rate (Pardoe and Stone 2006b). Acceptance probability distributions could also be trained off-line (Benisch et al. 2006). Another option is to try to model the decision function of the accepting entities, based on their decision histories, e.g., using Chebychev polynomials (Saha, Biswas, and Sen 2005).

Related work suggests some other aspects besides acceptance probability estimations to be relevant for product pricing. Current and future offers of competitors (outside options) could be considered (Li, Giampapa, and Sycara 2004). In (Sohn, Moon, and Seok 2009), outside options are considered as well. Here, a dynamic product pricing model is proposed in which the price change of the product itself as well as the relative price of competing products is quantified in a price elasticity. Using scenario analysis for distinguishing between various situations of price elasticity (i.e., market conditions), the optimal pricing policy is selected.

Market conditions can also be accounted for by using economic regimes (Ketter et al. 2006; 2007; 2009). A regime is a set of conditions, characterizing the state of a system or process. Regimes provide an intuitive way of conditioning behavior in different scenarios. In literature, several approaches to regime identification and prediction have been proposed in different contexts (Becker, Hurn, and Pavlov 2007; Hamilton 1989; Hathaway and Bezdek 1993; Mount, Ning, and Cai 2006). Regimes are useful in an economic context, because the ability of decision makers to correctly identify the current regime and predict the onset of a new regime is crucial in order to prevent over- or underreaction to market conditions (Massey and Wu 2005). Economic regimes can guide tactical (e.g., product pricing) and strategic sales decisions (e.g., product mix and production planning) (Ketter et al. 2006; 2007; 2009).

Finally, in case of known demand and uncertain supply, a responsive pricing policy, in which the retail price is determined after observing the realized supply, results in a higher expected profit than a pricing policy in which the realized supply is not taken into account (Tang and Yin 2007). Consequently, modeling expected or observed supply-side behavior in the product pricing process could contribute to profit maximization.

In the context of the TAC SCM game, several approaches to product pricing have been proposed. TacTex (Pardoe and Stone 2006b) predicts demand (using a Bayesian approach introduced by the DeepMaize team (Kiekintveld et al. 2004)) and offer acceptance (using linear regression) and adapts these offer acceptance estimations to its opponents' behavior. Another approach is to directly model the behavior of the competing agents and thus predict their offers (Kovalchuk and Fasli 2008). The SouthamptonSCM (He et al. 2005) agent uses fuzzy reasoning for daily price adaptation. For predicting whether a particular price for a particular product will be accepted by a customer, the CMieux agent (Benisch et al. 2006) uses probability distributions trained off-line. The Botticelli agent (Benisch et al. 2004) and PhantAgent (Stan, Stan, and Florea 2006) use simple heuristics for determining what to sell for what price.

A Product Pricing Model based on Offer Price Distributions

We propose to approximate acceptance probabilities in a scenario as simulated in the TAC SCM game by taking into account all offered prices, hereby implicitly modeling the decision making processes of all traders. Incorporating offer price distributions rather than individual offers into the framework can compensate for a drawback encountered in (Kovalchuk and Fasli 2008), where offer prices of individual competitors are predicted even though these competitors may not actually bid. When reasoning in terms of offer price distributions rather than individual offers, the offer price distributions are formed by all offers actually done rather than offers all competitors would make, should they actually bid.

Dynamic Pricing When customers only consider bids at or below their reservation price and always select the the cheapest offer, the distribution of realized order prices can be derived from the distribution of valid offer prices. Modeling order price distributions using offer price distributions is intuitive, as this captures the market dynamics and facilitates representation of relevant information on supply-side behavior. Now, let for product g on game day d a mean number of $\overline{n}_{dg}$ randomly sampled valid offer prices p_{dg} for each out of m_{dg} RFQs be identically and independently distributed in accordance with a distribution $f\left(p_{dg};\theta\right)$ and a cumulative distribution $F\left(p_{dg};\theta\right)$, with $0 < p_{dg} < u$ (i.e., prices are non-negative and have an upperbound u) and θ a vector of unknown parameters. For such a distribution, the cumulative distribution of the minimum valid offer prices (and thus the order prices) $\underline{p}_{dg}$ over all m_{dg} RFQs can be derived as (Kapadia, Chan, and Moyé 2005)

$$F_{\underline{p}}\left(\underline{p}_{dg};\theta\right) = \left(1 - \left(1 - F\left(\underline{p}_{dg};\theta\right)\right)^{\overline{n}_{dg}}\right)^{m_{dg}}, \quad 0 < \underline{p}_{dg} < u, \quad \overline{n}_{dg}, m_{dg} > 0. \tag{1}$$

The cumulative density of order prices yields the fraction of order prices realized at or below a specific value, which is similar to the probability that an order is placed with another trader offering a similar or better deal. Consequently, the reverse of this cumulative density approximates the probability for an agent to offer a better deal than other competitors. Hence, the reverse cumulative density of order prices associated with product g on game day d is an estimation of the probability that a customer will place an order o with an agent, given its offer price p_{dg}, $P\left(o|p_{dg}\right)$. Acceptance probabilities can therefore be estimated as

$$\begin{aligned} P\left(o|p_{dg}\right) &= \left(1 - F_{\underline{p}}\left(p_{dg};\theta\right)\right), \quad 0 < p_{dg} < u \\ &= \left(1 - \left(1 - \left(1 - F\left(p_{dg};\theta\right)\right)^{\overline{n}_{dg}}\right)^{m_{dg}}\right), \\ &\quad 0 < p_{dg} < u, \quad \overline{n}_{dg}, m_{dg} > 0. \end{aligned} \tag{2}$$

Equation (2) can be used to estimate the share of received orders with respect to the total number of RFQs for product g on game day d, generated by a price offered on all these

RFQs. Now, let q^*_{dg} be the sales quota (i.e., desired acceptance probability (Ketter et al. 2008)) for product g on day d, with m_{dg} associated RFQs, for each of which $\overline{n}_{dg}$ prices are offered. This implies that $P\left(o|p^*_{dg}\right)$ is required to be q^*_{dg}. Solving the equation to p^*_{dg} yields the optimal offer price p^*_{dg} expected to yield the desired quota. This way, products can be priced using estimations of offer price distributions.

Model Parameter Estimation When pricing products, the unknown parameters θ and $\overline{n}_{dg}$ in the acceptance probability approximation detailed in (2) must be estimated for a product g on game day d. When all data is available, $\overline{n}_{dg}$ can be determined by a counting process. Furthermore, θ can be estimated by maximum likelihood as follows. Assuming all prices in the sample of prices $\vec{p}_{dg}$ offered for all m_{dg} RFQs issued for product g on game day d to be identically and independently distributed in accordance with the offer price distribution $f\left(p_{dgri};\theta\right)$, the joint distribution of all valid offer prices can be derived as

$$f\left(\vec{p}_{dg};\theta\right) = \prod_{r=1}^{m_{dg}} \prod_{i=1}^{n_{dgr}} f\left(p_{dgri};\theta\right), \quad 0 < p_{dgri} < u, \quad (3)$$

where p_{dgri} is the ith of n_{dgr} prices offered on game day d for RFQ r for product g and parameters θ can be estimated by minimizing the negative log-likelihood function of these parameters for a sample of observed offer prices $\vec{p}_{dg}$ (e.g., using the Newton-Raphson method described in (Coleman and Li 1996)):

$$L\left(\theta;\vec{p}_{dg}\right) = \sum_{r=1}^{m_{dg}} \sum_{i=1}^{n_{dgr}} -\ln\left(f\left(p_{dgri};\theta\right)\right), \quad 0 < p_{dgri} < u. \quad (4)$$

However, data on offer prices is not available in TAC SCM, due to limited visibility of the real-time environment. Hence, the distribution parameters θ and $\overline{n}_{dg}$ can be estimated using a vector of real-time available information, $\vec{x}$. The relations between available information and distribution parameters, $h_\theta\left(\vec{x}\right)$ and $h_{\overline{n}_{dg}}\left(\vec{x}\right)$, can be modeled using Artificial Neural Networks (ANNs), yielding approximations of these relations: $\hat{h}_\theta\left(\vec{x}\right)$ and $\hat{h}_{\overline{n}_{dg}}\left(\vec{x}\right)$. An ANN is a mathematical model inspired by biological neural networks, which provides a general, practical method for learning real-valued, discrete-valued, and vector-valued functions over continuous and discrete-valued attributes from examples in order to facilitate regression or classification (Mitchell 1997). The model consists of interconnecting artificial neurons (nodes), ordered into an input layer, hidden layers, and an output layer.

Due to the ability of an ANN of capturing complex non-linear relations, which is a useful feature in case of learning functions whose general form is unknown in advance, parameter estimation using (4) can be replaced with such a model, albeit with different inputs (i.e., real-time available data). Representing the unknown relations between distribution parameters and data using ANNs also brings the attractive feature of fast evaluation of these (modeled) functions, which is crucial in case of real-time product pricing. Other advantages include robustness to noise in the training data (Mitchell 1997), the possibility to introduce adaptivity by adjusting the weights of each node's inputs on-the-fly using newly obtained examples (if any), and the fact that ANNs have proven to be useful for economic forecasts in various domains (Kovalchuk and Fasli 2008). Moreover, our experimental results regarding TAC SCM showed that ANNs better captured the relation between data and distribution parameters than for instance a linear regression model.

We propose to use a specific type of ANN for parameter estimation: a Radial Basis Function Network (RBFN). A RBFN is a two-layer ANN consisting of a hidden layer and an output layer. The activation function in each hidden unit h is a local function $K_h\left(d\left(\vec{x}_h,\vec{x}\right)\right)$, the output of which approaches 0 as $d\left(\vec{x}_h,\vec{x}\right)$ – the (typically Euclidian) distance between an instance characterized by a vector of features $\vec{x}$ and the center $\vec{x}_h$ – increases. The local functions in the hidden layer typically are Gaussians, centered at $\vec{x}_h$ with variance σ^2_h. The number of Gaussians H is subject to optimization and their centers can be determined by clustering the data, using for example the k-means algorithm (MacQueen 1967). The network's output for an instance $\vec{x}$, $\hat{h}\left(\vec{x}\right)$, is a linear combination of the activation units, weighted for their weights w_h, and a bias w_0 (Mitchell 1997):

$$\hat{h}\left(\vec{x}\right) = w_0 + \sum_{h=1}^{H} w_h K_h\left(d\left(\vec{x}_h,\vec{x}\right)\right), \quad (5)$$

$$K_h\left(d\left(\vec{x}_h,\vec{x}\right)\right) = e^{\frac{1}{2\sigma^2_h} d^2\left(\vec{x}_h,\vec{x}\right)}. \quad (6)$$

Hence, a RBFN is a global approximation $\hat{h}\left(\vec{x}\right)$ of a target function $h\left(\vec{x}\right)$, represented as a linear combination of local functions around their centers. Because RBFNs can be designed and trained in a fraction of the time it takes to train standard feed-forward back-propagation neural networks (Mitchell 1997), a RBFN would be a good approximator for distribution parameters.

Adaptive Regime-Based Product Pricing

In (Walsh et al. 2008), an English auction scenario is considered, in which bidders have independent private values, all originating from the same distribution. These private values result in bids up to the private values. The best (highest) bid wins. The distribution of the private values of the bidders is estimated using averaging and binary search techniques, combined with simulations. In this approach, adaptivity to market disruptions is realized by assuming changes in bidding (and thus market disruptions) to actually be a shift in the underlying value distribution. Therefore, the estimated distribution of private values is shifted accordingly.

The product pricing process we target is somewhat similar to the scenario described in (Walsh et al. 2008). In our case, traders bid on an RFQ and the best (lowest) bid wins. However, our RFQ bidding processes resemble (reverse) sealed bid, first price auctions rather than English auctions, as TAC SCM trading agents are not aware of bids of their competitors and the best (lowest) bid wins. Therefore, changes in bidding behavior of the competitors cannot be observed.

However, information on economic regimes might help in modeling changes in bidding behavior of manufacturers participating in a market, as realized order prices and hence the associated order probabilities tend to vary, depending on the economic regime (Ketter et al. 2006; 2007; 2009). Hence, changes in pricing behavior can be accounted for by incorporating regime information into the process of estimating order price distributions and the associated customer offer acceptance probabilities (Ketter et al. 2008). Therefore, we propose to dynamically model the relations between available data and price distributions using economic regimes. To this end, distribution parameters θ_k and $\overline{n}_{dgk}$ for product g on day d can be estimated per dominant regime k using RBFNs. When M dominant regimes are considered, this yields M separate price distribution estimations. The acceptance probabilities $P\left(o_k|p_{dgk}\right)$ derived from these distributions can be weighted with the associated regime probabilities for regime R_{dgk}, $P\left(R_{dgk}\right)$.

The weights in the RBFNs could be updated on-line, based on new data. However, when new training samples cannot be presented to the networks (due to limited visibility of market characteristics), daily acceptance probability estimations done using the RBFNs can be adjusted by multiplying the acceptance probabilities by a factor representing the ratio between the number of actually received orders and the number of predicted orders, as proposed in (Pardoe and Stone 2006a). If more orders have been received than predicted, the acceptance probability is larger than expected, to an extent equal to the ratio between received and predicted number of orders. If less orders have been received than predicted, the acceptance probability should be adjusted downwards. This ratio, which can also be referred to as a residual error term ϵ, enables market responses to be fed back to the model, as this ratio can be updated in real-time. A smoothed error term $\widetilde{\epsilon}$ can be used in order to prevent over- or under-compensation.

For dominant regime k, the probability that a customer accepts an offer of a manufacturer and hence places an order o_k with this manufacturer, given price p_{dgk} for product g on game day d, $P\left(o_k|p_{dgk}\right)$, ranges from 0 to 1. Multiplying this customer offer acceptance probability with the suggested ratio $\widetilde{\epsilon}_{(d-1)gk}$ (which depends on regime k and has been updated using performance information up until day $d-1$) yields a corrected customer offer acceptance probability $P\left(o_k|p_{dgk}\right)'$ in the range $\left[0, \widetilde{\epsilon}_{(d-1)gk}\right]$. This implies that no suitable price can be found for $q^*_{dg} \geq \widetilde{\epsilon}_{(d-1)gk}$, which is an undesirable feature when $\widetilde{\epsilon}_{(d-1)gk} < 1$. However, when the corrected customer offer acceptance probability $P\left(o_k|p_{dgk}\right)'$ is defined as

$$\begin{aligned}
&P\left(o_k|p_{dgk}\right)' \\
&= P\left(o_k|p_{dgk}\right)^{\widetilde{\epsilon}_{(d-1)gk}}, \quad 0 < p_{dgk} < u, \quad \widetilde{\epsilon}_{(d-1)gk} > 0, \\
&= \left(1 - \left(1 - \left(1 - F\left(p_{dgk}; \theta_k\right)\right)^{\overline{n}_{dgk}}\right)^{m_{dg}}\right)^{\widetilde{\epsilon}_{(d-1)gk}}, \\
&\quad 0 < p_{dgk} < u, \quad \overline{n}_{dgk}, m_{dg}, \widetilde{\epsilon}_{(d-1)gk} > 0, \qquad (7)
\end{aligned}$$

customer offer acceptance probabilities continue to range from 0 to 1 for $0 < p_{dgk} < u$ after correction.

Using (7), the corrected offer price $p^{*'}_{dg}$ expected to yield the desired sales quota q^*_{dg} for product g on day d can be defined for each dominant regime k, by requiring $P\left(o_k|p^{*'}_{dgk}\right)'$ for that product on that day to be q^*_{dg}. Solving the equation to $p^{*'}_{dgk}$ yields the optimal corrected offer price $p^{*'}_{dgk}$ expected to yield the desired quota q^*_{dg} under dominant regime k. When these corrected prices are then weighted for their associated regime probabilities $P\left(R_{dgk}\right)$, the corrected price $p^{*'}_{dg}$ expected to yield the required quota can be obtained.

The error term considered in (7) should be assigned values such that under each dominant regime k, the expected customer offer acceptance probabilities $P\left(o_k|p^{*'}_{(d-1)g}\right)$ associated with an offer price $p^{*'}_{(d-1)g}$ are corrected by the unsmoothed exponential error terms to the proportion of actually received number of orders $q_{(d-1)g}$. The found error terms can subsequently be smoothed. Hence, offer price and customer response should be related as shown in (8).

$$\begin{aligned}
q_{(d-1)g} &= P\left(o_k|p^{*'}_{(d-1)g}\right)^{\epsilon_{(d-1)gk}}, \quad 0 < q_{(d-1)g} < 1, \\
&\quad 0 < p^{*'}_{(d-1)g} < u, \quad \epsilon_{(d-1)gk} > 0. \qquad (8)
\end{aligned}$$

The error terms can be smoothed using double exponential smoothing (Brown, Meyer, and D'Esopo 1961), where the smoothing factor β is weighted for the associated regime probabilities in order for errors only to be attributed to the models responsible for these errors. Smoothing is done by linearly combining two components (see (12)), the first of which (defined in (10)) is a linear combination of the latest error (see (9)) and the previous first component. The second component (defined in (11)) is a linear combination of the first component and the previous second component.

$$\begin{aligned}
\epsilon_{(d-1)gk} &= \frac{\ln\left(q_{(d-1)g}\right)}{\ln\left(P\left(o_k|p^{*'}_{(d-1)g}\right)\right)}, \quad 0 < q_{(d-1)g} < 1, \\
&\quad 0 < P\left(o_k|p^{*'}_{(d-1)g}\right) < 1, \qquad (9) \\
\widetilde{\epsilon}'_{(d-1)gk} &= \beta P\left(R_{(d-1)gk}\right)\epsilon_{(d-1)gk} + \\
&\quad \left(1 - \left(\beta P\left(R_{(d-1)gk}\right)\right)\right)\widetilde{\epsilon}'_{(d-2)gk}, \qquad (10) \\
\widetilde{\epsilon}''_{(d-1)gk} &= \beta P\left(R_{(d-1)gk}\right)\widetilde{\epsilon}'_{(d-1)gk} + \\
&\quad \left(1 - \left(\beta P\left(R_{(d-1)gk}\right)\right)\right)\widetilde{\epsilon}''_{(d-2)gk}, \qquad (11) \\
\widetilde{\epsilon}_{(d-1)gk} &= 2\widetilde{\epsilon}'_{(d-1)gk} - \epsilon''_{(d-1)gk}. \qquad (12)
\end{aligned}$$

So far, the proposed framework assumes the offer prices to be distributed in accordance with a distribution the type and parameters of which have not been defined yet. For this purpose, we propose a log-logistic distribution, as this distribution covers a variety of shapes depending on the parameters α and γ and has the attractive feature that an analytical closed form expression exists for the cumulative density function. Moreover, we determined that this distribution sufficiently describes the data in the TAC SCM game in over 60% of the analyzed price samples from historical game data[1], according to the Kolmogorov-Smirnov test (Massey 1951) (when requiring the p-value to be over 0.05). The remaining samples could not sufficiently be described using a simple parametric distribution such as the log-logistic distribution due to the complex form of their true densities. The log-logistic distribution $f(p;\alpha,\gamma)$ and its cumulative form $F(p;\alpha,\gamma)$ (Mood, Graybill, and Boes 1974), truncated such that the distribution is defined on the domain $0 < p < u$ and reparameterized such that α represents the median and γ quantifies the distribution tightness, can be described as

$$f(p;\alpha,\gamma) = \frac{\left(\alpha^{-\gamma} - u^{-\gamma}\right)\gamma p^{-\gamma-1}}{\left(\alpha^{-\gamma} - 2u^{-\gamma} + p^{-\gamma}\right)^2}, \quad 0 < p < u, \quad \alpha, \gamma > 0, \tag{13}$$

$$F(p;\alpha,\gamma) = \frac{\alpha^{-\gamma} - u^{-\gamma}}{\alpha^{-\gamma} - 2u^{-\gamma} + p^{-\gamma}}, \quad 0 < p < u, \quad \alpha, \gamma > 0. \tag{14}$$

When $F(p_{dgk};\theta_k)$ in (7) is substituted for (14), the corrected price $p^{*'}_{dg}$ expected to yield a required sales quota q^*_{dg} can be obtained as shown in (15) through (17). Here, let $1 \leq k \leq M$, with the number of considered regimes $M = 5$ (Ketter et al. 2006; 2007; 2009).

$$q^{*'}_{dg} = \sqrt[\overline{n}_{dgk}]{1 - \sqrt[m_{dg}]{1 - \sqrt[\widetilde{\epsilon}_{(d-1)gk}]{q^*_{dg}}}}, \quad 0 < q^*_{dg} < 1, \quad m_{dg}, \overline{n}_{dgk}, \widetilde{\epsilon}_{(d-1)g} > 0, \tag{15}$$

$$p^{*'}_{dgk} = \left(\frac{u^{-\gamma_k}\left(\alpha_k^{-\gamma_k}\left(u^{\gamma_k} - 2\alpha_k^{\gamma_k}\right)q^{*'}_{dg} + 1\right)}{1 - q^{*'}_{dg}}\right)^{-\frac{1}{\gamma_k}}, \quad \alpha_k, \gamma_k > 0, \tag{16}$$

$$p^{*'}_{dg} = \sum_{k=1}^{5} P(R_{dgk})\, p^{*'}_{dgk}, \quad 0 < p^{*'}_{dgk} < u. \tag{17}$$

We now have a product pricing approach, which assumes a double-bounded log-logistic distribution to be underlying offer prices, the parameters of which can be estimated real-time using RBFNs, based on available information. This approach can adapt to market disruptions, characterized using economic regimes, as product prices are determined per dominant regime using (15) and (16) and weighted for their associated regime probabilities in (17). The relations between price distributions and available information are thus dynamically modeled, depending on economic regimes.

Structural errors in the product pricing process are accounted for by feeding back market responses to placed offers using an error term, designed to transform the estimated acceptance probability function into a function better approximating the true acceptance probability. This error term is corrected using daily observations of expected and observed acceptance probabilities – double exponentially smoothed with a smoothing factor weighted for the associated regime probabilities – using (9) through (12). This feedback process enables the product pricing model to adapt to the true customer offer acceptance probabilities.

[1] TAC SCM 2007 Semi-Finals and Finals (9323–9327tac5 and 7308–7312tac3) (SICS AB 2004 2009) and TAC SCM 2008 Semi-Finals and Finals (763–768tac02 and 794–799tac01) (University of Minnesota 2003 2009) game data.

Performance in TAC SCM

The limited visibility of the TAC SCM market environment advocates the need for effective modeling of the market environment, such that decision making processes are adequately supported. As economic regimes have been shown to be useful in effectively describing market conditions in TAC SCM (Ketter et al. 2006; 2007; 2009), dynamically differentiating pricing strategies using these economic regimes would intuitively contribute to the performance of a manufacturer within the game.

Therefore, the final framework is evaluated by implementing the approach in the MinneTAC agent (Collins, Ketter, and Gini 2009) for TAC SCM. To this end, product pricing should be done using (9) through (12) and (15) through (17). The α_k, γ_k, and $\overline{n}_{dgk}$ parameters for product g on game day d for dominant regime k are to be estimated using RBFNs.

Implementation in the MinneTAC Agent The sales decisions made by the MinneTAC agent originate from price predictions based on microeconomic conditions, characterized for each individual market segment: economic regimes are identified and predicted. These economic regimes can be extreme scarcity, scarcity, a balanced situation, oversupply, and extreme oversupply. On game day d, the regime for good g is identified using regime probabilities. The regime having the highest probability, given the estimated normalized mean price of that day is the current dominant regime. This price estimate is a smoothed normalized mid-range price, $\widetilde{\mathrm{np}}_{dg}$, which is the average of the double exponentially smoothed normalized minimum and maximum price. To this end, prices are normalized by expressing these prices in terms of their associated production costs, such that these normalized prices np_{dg} range from 0 to 1.25.

A product-level price density function has been modeled on historical normalized order price data using a Gaussian Mixture Model (GMM) (Titterington, Smith, and Makov 1985) with a sufficient number of Gaussians (25 at the moment), reflecting a balance between prediction accuracy and computational overhead. Clustering these price distributions over time periods (using the k-means algorithm) has yielded distinguishable statistical patterns: economic regimes. In TAC SCM, price information is only available up until the preceding game day. Hence, the MinneTAC agent estimates the mean price of day d using exponential smoothing prediction of $\widetilde{\mathrm{np}}_{dg}$ and then returns the regime probabilities for day d. When such an estimate is supplied to the model, the individual Gaussians in the model are activated to a certain extent, thus generating an expected price distribution. Subsequently normalizing all clusters' price densities enables determination of regime probabilities. Future regime probabilities are determined by using Markov prediction and Markov correction-prediction processes (for short-term and long-term decision making, respectively) (Ketter et al. 2008). When current or future regime probabilities have been determined, products are priced using acceptance probabilities of these prices, such that a sales quota is fulfilled.

In the sales process configuration used as benchmark in our research, price trends are estimated by the regime model. The median price of a product is estimated using a price-following approach implementing a double exponential smoother. These trends, combined with the estimated median, are used in the allocation process, where sales quotas are generated based on – among other things – these price predictions. The curve representing the probability of acceptance is approximated using the estimated median price and the curve's slope in that median, estimated using exponentially smoothed prices. This acceptance probability is used for determining the price to be offered in order for the sales agent to sell its desired quota.

In order to compensate for the uncertainty in generated predictions, interval randomization is applied to offer prices, which adds a slight variability to these prices. The estimated median is corrected using feedback derived from the desired customer offer acceptance probability and the associated true acceptance probability observed the next day. A major drawback here is the assumption that customer feedback is in response to the optimized offer price, whereas this feedback is in response to a price randomized in an interval around this price.

The price distributions estimated by the GMM are updated on-line, but they do not account for factors other than a mean price estimate and lack full adaptivity. In an attempt to improve the product pricing process by combining regime information with real-time available data, we replace the sales model of the benchmark with a system designed for Product Pricing using Adaptive Real-time Regime-based Probability of Acceptance Estimations: PPARRPAE. The algorithm (see Algorithm 1) involves parameter estimation using RBFNs and subsequently pricing products using (15) through (17) (with $u = 1.25$). In this process, an error term is considered, following (9) through (12) (with $\beta = 0.5$, as determined by a hill-climbing procedure).

The idea is to build an adapter, which combines characteristics of price distribution estimated by the regime model with available information. Using the RBFNs, the adapter transforms available information into a parameterized acceptance probability function per dominant regime and assigns weights to these functions, equal to their associated regime probabilities. This adapted distribution can subsequently be used in the product pricing process. Given a quota for a product (specified by the allocation component), the product pricing component uses the adapter to compute the price expected to yield this quota per dominant regime and weights the suggested prices for their associated regime probabilities. The optimized price is then offered on all selected RFQs for the considered product. The market responses to these offers are fed back to the adapter. In order for this information not to be biased, interval randomization is not applied to the optimized price, as opposed to the benchmark approach.

```
foreach d in days do
    foreach g in products do
        // Update error using last
        // feedback, using (9) through (12)
        error = updateError(getFdback(d − 1, g));
        // Retrieve data from regime model
        regProbs = getRegProbs(d, g);
        regPriceDistr = getRegPriceDistr(d, g);
        trends = getTrends(d, g);
        // Estimate parameters using RBFNs
        priceDistr = estParams(regPriceDistr,
        getData(d, g));
        // Determine median price using (15)
        // through (17)
        median = priceForProb(0.5, priceDistr,
        error, regProbs);
        // Retrieve allocated quota
        quota = getQuota(d, g, median, trends);
        // Determine price expected to
        // yield quota using (15) through (17)
        price = priceForProb(quota, priceDistr,
        error, regProbs);
        // Bid price on selected RFQs
        priceProduct(d, g, price);
    end
end
```

Algorithm 1: The PPARRPAE approach.

The allocation model partially bases its decisions on price predictions, which consist of an estimate of the median price of the considered game day and trends representing expected future deviations from this median. In the benchmark sales model, the regime model provides trends, whereas the median is estimated using a price-follower approach. This price-following component is also used for acceptance probability estimation and can thus be updated using market responses. Since in the proposed approach, market responses are not related to the price-following median, but are fed back to the adapter, the median price prediction should be provided to the allocation component by the adapter.

Radial Basis Function Network Training For each dominant regime k, a RBFN needs to be trained for estimating α_k, γ_k, and $\overline{n}_{dgk}$ for product g on game day d. Therefore, training and test datasets[2] must be split into datasets per dominant regime. These dominant regimes are identified by the current regime model. We attempt to adapt regime-based price distributions done using the GMM implemented in the MinneTAC agent in order for them to be useful in the daily product pricing process. Hence, these regime-based price distributions should be used as RBFN inputs. For now, let these distributions be described by their 10th, 50th, and 90th percentile and the spread of these percentiles. The RBFNs should adapt these distributions using on-line available data.

[2]TAC SCM 2007 Semi-Finals and Finals (9321–9328tac5 and 7306–7313tac3) (SICS AB 2004 2009) and TAC SCM 2008 Semi-Finals and Finals (761–769tac02 and 792–800tac01) (University of Minnesota 2003 2009) game data. The first two games and the last game per server form the test set, the rest forms the training set.

TAC SCM price distributions tend to differ per product type, so product types can indicate price distribution characteristics. Offered prices can also be related to game days. E.g., in the game's start-up phase, prices are more likely to be relatively high due to product scarcity. Another indicator for a product's offer price distribution can be the number of RFQs for that product, as the number of simultaneously run similar auctions affects the generated revenue due to the auctions' (partial) substitutivity (Walsh et al. 2008). In TAC SCM, RFQs for a product type can be (partial) substitutes to bidding agents, as their production capacities are limited. RFQ characteristics can indicate the pricing behavior they generate too, so mean and standard deviation of requested quantities, leadtimes, and reservation prices can be considered. Prices realized on the preceding day are useful too (Kovalchuk and Fasli 2008). In-game, only a product's minimum and maximum order prices (and their midrange and spread) realized on the preceding day are visible. These prices can be double exponentially smoothed, yielding a good estimate of the associated mean price (Ketter et al. 2006; 2007; 2009).

The RBFNs should hence be trained to adapt regime-based GMM price distribution estimates using data on product type, game day, RFQ characteristics, and observable prices. Using historical game data, target parameter values can be determined by a counting process (for $\overline{n}_{dgk}$) and by fitting distributions using (4) and (13) (for α_k and γ_k). A typical training dataset thus generated contains over 15,000 samples, a typical test dataset over 8,000. The performance of RBFNs trained on the training set can be evaluated on the test set, as the latter set is sufficiently large and representative (Mitchell 1997). The thus found optimal values for γ_k tend to be distributed on an exponential scale; the increment in γ_k needed to tighten the distribution increases as the distribution gets tighter. E.g., a distribution with a γ_k value of 2 is much more different from one with a γ_k value of 5 than a distribution with a γ_k value of 200 is from one with a γ_k value of 500, when all other parameters are fixed. Hence, as the required accuracy decreases for an increasing γ_k, the networks are trained to predict the natural logarithm of γ_k.

Using Weka (Witten and Frank 2005), the RBFNs can be trained relatively easily. The results can be saved as serialized Java objects, which can be used in Java software like MinneTAC. Another advantage of Weka is the availability of other model types. As most other models implemented in Weka are classification rather than regression models and preliminary tests indicate the preferability of the RBFN implementation (`RBFNetwork`) over other regression models, we use `RBFNetwork`. One drawback is that `RBFNetwork` can only have one output. Hence, we train a RBFN per dominant regime per parameter. We fix the random seed `RBFNetwork` uses in the clustering process used to determine the centers of the Gaussians in the networks. One can also specify a so-called ridge value, which indicates how much the regression error in estimating model parameters may diverge from the least squares measure. For all networks, this value is left at its default value, 1E-08. Other parameters are the number of clusters and the minimum standard deviation of these clusters.

Parameter	Regime	Clusters	MinStdev	RMSD
α_k	1	25	15	0.0448
α_k	2	50	10	0.0346
α_k	3	100	5	0.0366
α_k	4	50	5	0.0386
α_k	5	300	5	0.0400
$\ln(\gamma_k)$	1	100	15	0.7713
$\ln(\gamma_k)$	2	150	5	0.6903
$\ln(\gamma_k)$	3	150	5	0.6481
$\ln(\gamma_k)$	4	200	5	0.6370
$\ln(\gamma_k)$	5	150	2	0.6732
$\overline{n}_{dgk}$	1	50	15	1.0036
$\overline{n}_{dgk}$	2	25	5	1.0773
$\overline{n}_{dgk}$	3	200	5	0.9974
$\overline{n}_{dgk}$	4	300	2	0.9395
$\overline{n}_{dgk}$	5	100	5	0.8090

Table 1: Optimized configuration of number of clusters and minimum standard deviation of clusters for RBFNs estimating distribution parameters, along with the RMSD of parameter values predicted by the models from their target values.

The configurations of the latter two parameters can be determined by systematically evaluating all combinations of different values. The configurations yielding the lowest root mean squared deviation (RMSD) on the test set are selected (Mitchell 1997). The RMSD can be defined as

$$\text{RMSD} = \sqrt{\frac{\sum_{\omega=1}^{\Omega}\left(\hat{x}_\omega - x_\omega\right)^2}{\Omega}}, \tag{18}$$

where $\hat{x}_\omega$ is a predicted observation in a set of Ω observations and x_ω is the associated actual value.

The optimal number of clusters could be anything between relatively small and rather large. Using too many clusters would cause the model to not generalize very well. Hence, taking into account the size of the dataset, the set of number of clusters considered is $\{25, 50, 100, 150, 200, 300\}$ and standard deviations in the set $\{1, 2, 5, 10, 15\}$ are considered. Apparently, α_k (ranging between 0 and 1.25) can be estimated relatively well, whereas $\overline{n}_{dgk}$ (ranging between 0 and 6) and $\ln(\gamma_k)$ (roughly ranging between -6 and 6) cannot (see Table 1).

Performance Evaluation By running and analyzing a number of games, the performance of the proposed PPAR-RPAE system can be compared with the performance of the benchmark. In this experimental set-up, games are in accordance with the 2006 TAC SCM game specifications (Collins et al. 2005). The randomness incorporated in the game (e.g., in customer demand) is an inconvenient characteristic for a testing environment in which two approaches are to be compared, as this randomness in market conditions implies that many experiments should be run in order to obtain results with any statistical significance. The issue of randomness in the testing environment is tackled by a controlled TAC SCM game server, in which random seeds – which are used for generating market conditions – can be controlled. Uncontrolled stochastic behavior of participating trading agents does not have a significant impact on the overall agent profit levels and most significant profit differences between agents can already be detected in approximately 40 games (Sodomka, Collins, and Gini 2007).

The performance of the PPARRPAE system can hence be evaluated in 40 experiment sets on a controlled server. Each experiment set consists of a paired evaluation of the performance of the benchmark and the PPARRPAE system under equal market characteristics. We let the competitors be the default competitors that come with the TAC SCM game. These competitors use a make-to-order strategy. In each evaluation, the final bank account balance can be considered, as well as the sales performance. To this end, the mean and standard deviation of account balances over all games can be computed. The number of obtained orders should be considered in the analysis as well. The number of times the agent proceeds to actually bidding on RFQs, given an acceptance probability estimate, can also be analyzed.

Performance differences should also be assessed with respect to their statistical relevance. This can be done with a paired Student's t-test (Goulden 1956), which tests whether the population means of two samples are equal (this null hypothesis is rejected at a significance level below 0.05). However, this statistic assumes the observations to be distributed in accordance with a normal distribution and we do not know whether this is a realistic assumption. Therefore, we also assess statistical relevance of observed performance differences using the paired, two-sided Wilcoxon signed-rank test (Gibbons 1986; Hollander and Wolfe 2000). This is a non-parametric test, which tests the hypothesis that the differences between paired observations are symmetrically distributed around a median equal to 0. If this null hypothesis is rejected (at a significance level below 0.05), the compared sets of samples can be assumed to be significantly different. This test is suitable in this experimental setup, as the distribution of the values to be compared is unknown.

Over all experiments, PPARRPAE turns out to outperform the benchmark with respect to final bank account balance and the number of obtained orders (see Tables 2 and 3). Generally, when using PPARRPAE, final balances significantly increase with about 160% (with a p-value of 0.0000 for both the paired t-test and the Wilcoxon test) and the number of orders significantly increases with over 50% (with a p-value of 0.0000 for the paired t-test as well as for the Wilcoxon test) with respect to the benchmark.

The increase in number of obtained orders can be explained by the significant increase of the number of usable acceptance probability estimations with approximately 80% (with p-values of 0.0000 for both tests). However, final account balance increase does not appear to be fully explained by an increase in obtained orders. In some experiments, a small increase (or even a decrease) in the number of obtained orders still results in doubled profits. This indicates that orders are better priced. This could be caused by prices of obtained orders to be closer to second-lowest prices, instead of being significantly lower, which results in a reduced margin between customers' reserve prices and realized order prices. Hence, using the PPARRPAE approach improves the quality of acceptance probability estimations and consequently results in better bid efficiency.

	Benchmark		PPARRPAE	
Agent	Mean	Stdev	Mean	Stdev
MinneTAC	19.2614	12.4207	49.3933	2.7053
Make-to-order-1	12.9194	3.2799	14.0436	2.9310
Make-to-order-2	13.0250	3.3152	14.1313	3.1668
Make-to-order-3	12.7687	3.3184	14.1034	2.7711
Make-to-order-4	12.8552	3.4148	14.3529	2.9034
Make-to-order-5	13.0803	3.2224	14.2307	2.9874

Table 2: Mean and standard deviation of final bank account balance per agent (in millions) over all experiments.

	Benchmark		PPARRPAE	
Agent	Mean	Stdev	Mean	Stdev
MinneTAC	3.0865	1.0178	4.6474	0.4507
Make-to-order-1	3.2615	0.3367	3.0498	0.3129
Make-to-order-2	3.2615	0.3291	3.0571	0.3448
Make-to-order-3	3.2452	0.3386	3.0364	0.3267
Make-to-order-4	3.2198	0.3292	3.0387	0.3172
Make-to-order-5	3.2504	0.3480	3.0437	0.3301

Table 3: Mean and standard deviation of number of obtained orders per agent (in thousands) over all experiments.

Conclusions and Future Work

When product pricing strategies are linked to price distribution estimations, taking into account real-time available information, the relation between available data and price distribution parameters can be dynamically modeled using economic regimes (characterizing market conditions) and error terms (accounting for customer feedback). Thus, in a constrained environment like TAC SCM, economic regime estimations turn out to contribute to profit maximization when they are used to differentiate product pricing strategies. Test results indicate that this novel approach significantly improves the performance of a price-following product pricing approach; more orders are obtained against higher prices and profits more than double. Much gain is obtained from using proper statistical methods, combined with effective real-time parameter estimation.

Even though the performance of the proposed model already is very promising, some aspects still require more research. First of all, the type and parameterization of models for real-time price distribution and acceptance probability approximation could be further improved. Other possible predictors for acceptance probabilities could be considered as well. Procurement information might be a good candidate here, as costs associated with specific orders could influence the price, depending on the cost allocation applied in the participating trading agents. Our current RBFNs mostly use information on RFQ characteristics as inputs in order to account for possibly dependent bidding behavior caused by (partial) substitutivity of RFQs per product type. Future research could therefore also include moving compensation for possible dependencies from the parameter estimation phase to the underlying model. Another option for future research could be trying to use the improved acceptance probability estimations in the allocation or RFQ selection process.

Finally, our approach of product pricing using adaptive regime-based acceptance probability estimations could be challenged in a situation with very tough competition. Currently, we are testing our approach against world's leading TAC SCM agents. If the MinneTAC agent could deal with those agents as with the agents considered in this research, MinneTAC would be more competitive than ever.

References

Becker, R.; Hurn, S.; and Pavlov, V. 2007. Modelling Spikes in Electricity Prices. *The Economic Record* 82(263):371–382.

Benisch, M.; Greenwald, A.; Grypari, I.; Lederman, R.; Naroditskiy, V.; and Tschantz, M. 2004. Botticelli: A Supply Chain Management Agent Designed to Optimize under Uncertainty. *ACM SIGecom Exchanges* 4(3):29–37.

Benisch, M.; Sardinha, A.; Andrews, J.; and Sadeh, N. 2006. CMieux: Adaptive Strategies for Competitive Supply Chain Trading. *ACM SIGecom Exchanges* 6(1):1–10.

Brown, R. G.; Meyer, R. F.; and D'Esopo, D. A. 1961. The Fundamental Theorem of Exponential Smoothing. *Operations Research* 9(5):673–687.

Coleman, T., and Li, Y. 1996. An Interior Trust Region Approach for Nonlinear Minimization Subject to Bounds. *SIAM Journal on Optimization* 6(2):418–445.

Collins, J.; Arunachalam, R.; Sadeh, N.; Eriksson, J.; Finne, N.; and Janson, S. 2005. The Supply Chain Management Game for the 2006 Trading Agent Competition. Technical report, Carnegie Mellon University.

Collins, J.; Ketter, W.; and Gini, M. 2009. Flexible decision control in an autonomous trading agent. *Electronic Commerce Research and Applications*.

Dasgupta, P., and Hashimoto, Y. 2004. Multi-Attribute Dynamic Pricing for Online Markets Using Intelligent Agents. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, 277–284. Washington, DC, USA: IEEE Computer Society.

Ghiani, G.; Laporte, G.; and Musmanno, R. 2004. *Introduction to Logistics Systems Planning and Control*. John Wiley & Sons, Ltd.

Gibbons, J. D. 1986. Nonparametric Statistical Inference. *Technometrics* 28(3):275.

Goulden, C. H. 1956. *Methods of Statistical Analysis*. John Wiley & Sons, Ltd, 2nd edition.

Hamilton, J. D. 1989. A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* 57(2):357–384.

Hathaway, R., and Bezdek, J. 1993. Switching Regression Models and Fuzzy Clustering. *IEEE Transactions on Fuzzy Systems* 1(3):195–204.

He, M.; Rogers, A.; David, E.; and Jennings, N. R. 2005. Designing and Evaluating an Adaptive Trading Agent for Supply Chain Management Applications. In Poutré, H. L.; Sadeh, N.; and Janson, S., eds., *IJCAI Workshop on Trading Agent Design and Analysis*, 35–42. Springer.

Hollander, M., and Wolfe, D. A. 2000. Nonparametric Statistical Methods. *Journal of the American Statistical Association* 95(449):333.

Kapadia, A. S.; Chan, W.; and Moyé, L. A. 2005. *Mathematical Statistics with Applications*. CRC Press.

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

Ketter, W.; Collins, J.; Gini, M.; Gupta, A.; and Schrater, P. 2007. A Predictive Empirical Model for Pricing and Resource Allocation Decisions. In *Proceedings of the Ninth International Conference on Electronic Commerce (ICEC 2007)*, 449–458.

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

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

Kiekintveld, C.; Wellman, M. P.; Singh, S. P.; Estelle, J.; Vorobeychik, Y.; Soni, V.; and Rudary, M. R. 2004. Distributed Feedback Control for Decision Making on Supply Chains. In *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, 384–392.

Kovalchuk, Y., and Fasli, M. 2008. Adaptive Strategies for Predicting Bidding Prices in Supply Chain Management. In *Proceedings of the Tenth International Conference on Electronic Commerce (ICEC 2008)*.

Li, C.; Giampapa, J.; and Sycara, K. 2004. Bilateral Negotiation Decisions with Uncertain Dynamic Outside Options. In *First IEEE International Workshop on Electronic Contracting*, 54–61. Washington, DC, USA: IEEE Computer Society.

MacQueen, J. 1967. Some Methods of Classification and Analysis of Multivariate Observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 281–297.

Massey, C., and Wu, G. 2005. Detecting Regime Shifts: The Causes of Under- and Overreaction. *Management Science* 51(6):932–947.

Massey, F. 1951. The Kolmogorov-Smirnov Test for Goodness of Fit. *Journal of the American Statistical Association* 46(253):68–78.

Mitchell, T. M. 1997. *Machine Learning*. McGraw-Hill Series in Computer Science. McGraw-Hill.

Mood, A.; Graybill, F.; and Boes, D. 1974. *Introduction to the theory of statistics*. McGraw-Hill.

Mount, T. D.; Ning, Y.; and Cai, X. 2006. Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters. *Energy Economics* 28(1):62–80.

Pardoe, D., and Stone, P. 2006a. Predictive Planning for Supply Chain Management. In *Proceedings of the Sixteenth International Conference on Automated Planning and Scheduling (ICAPS 2006)*, 21–30. Cumbria, UK: AAAI Press.

Pardoe, D., and Stone, P. 2006b. TacTex-05: A Champion Supply Chain Management Agent. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, 1489–94.

Saha, S.; Biswas, A.; and Sen, S. 2005. Modeling Opponent Decision in Repeated One-Shot Negotiations. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, 397–403. New York, USA: ACM.

SICS AB. 2004–2009. Trading Agent Competition (SICS). Website. Available online, `http://www.sics.se/tac`; last visited February 23, 2009.

Sodomka, E.; Collins, J.; and Gini, M. L. 2007. Efficient Statistical Methods for Evaluating Trading Agent Performance. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, 770–775. Vancouver, British Columbia, Canada: AAAI Press.

Sohn, S. Y.; Moon, T. H.; and Seok, K. J. 2009. Optimal Pricing for Mobile Manufacturers in Competitive Market Using Genetic Algorithm. *Expert Systems with Applications* 36(2):3448–3453.

Stan, M.; Stan, B.; and Florea, A. M. 2006. A Dynamic Strategy Agent for Supply Chain Management. In *Proceedings of the Eighth International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*, 227–232. Washington, DC, USA: IEEE Computer Society.

Tang, C. S., and Yin, R. 2007. Responsive Pricing under Supply Uncertainty. *European Journal of Operational Research* 182(1):239–255.

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

University of Minnesota. 2003–2009. Trading Agent Competition (UMN). Website. Available online, `http://tac.cs.umn.edu/`; last visited February 23, 2009.

Walsh, W. E.; Parkes, D. C.; Sandholm, T.; and Boutilier, C. 2008. Computing Reserve Prices and Identifying the Value Distribution in Real-world Auctions with Market Disruptions. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, 1499–1502.

Witten, I. H., and Frank, E. 2005. *Data Mining: Practical machine learning tools and techniques*. San Fransisco: Morgan Kaufmann, 2nd edition.