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

Wolf Ketter

RSM Erasmus University, The Netherlands
wketter@rsm.nl – www.ketter.ws

Man-Machine Interaction Group
Delft University of Technology, May 28th 2008

Work done with: John Collins, Maria Gini, Alok Gupta, and Paul Schrater

Outline

Motivation

- Economic theory suggests that economic environments exhibit 3 dominant market patterns: scarcity, balanced, and over-supply.
- We call these distinguishable conditions *economic regimes*.
- The long term objective of our work is to show how knowledge of current and anticipated regimes can enable an agent to make better operational and strategic decisions.

Relationship: Prices, Order Probability, and Regimes

Scarcity:

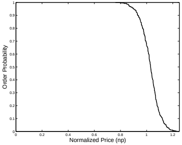

Balanced:

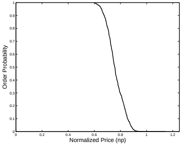

Over-supply:

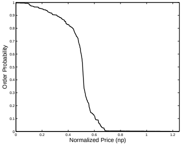

Reverse cumulative density function represents probability of order.

Experimental:

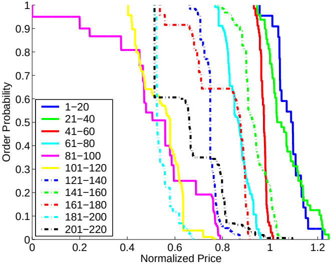

Application Areas (1)

- Identification of economic regimes:
  - Strategical decision making
  - Tactical decision making
- Price and price trend forecasting.
- Forecasting of economic regimes shifts:
  - Whole seller (e.g. book store).
  - Production plant (e.g. Daimler).
- Automated supply-chain management, e.g.,
  - i2
  - SAP

Application Areas (2)

- The approach we propose works in any market:
  - Computational process is completely data driven.
  - No classification of the market structure (monopoly vs competitive, etc) is needed.

TAC SCM
Overview

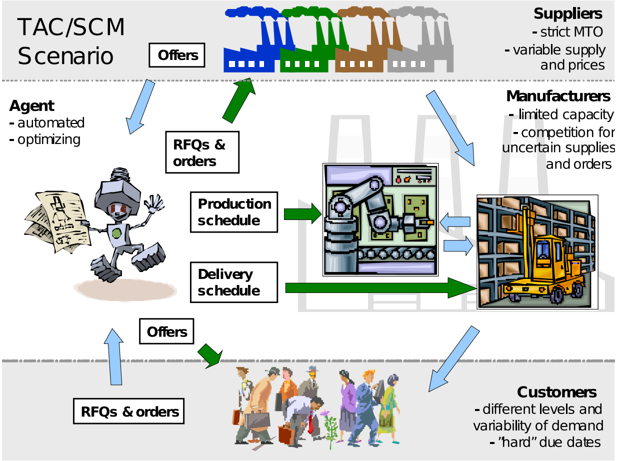

TAC SCM - Scenario

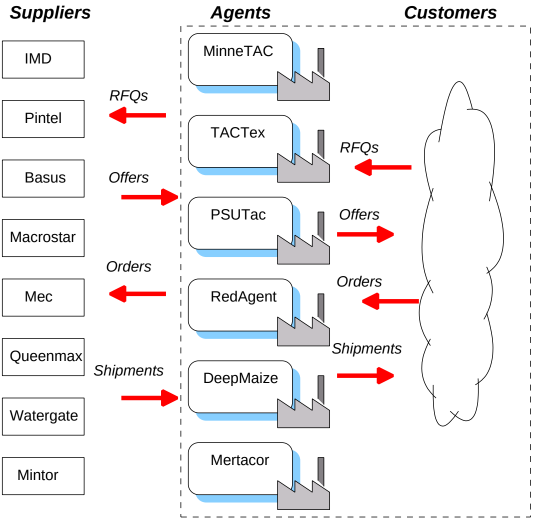

Use Regime Prediction For Sales Strategies

1. Allocation (Strategic Decision):
    - Allocating parts and production capacity to most profitable computers.
    - Allocating computers to current vs future sales.
2. Pricing (Tactical Decision):
    - Find the best prices to move the desired inventory.

Pricing Chain

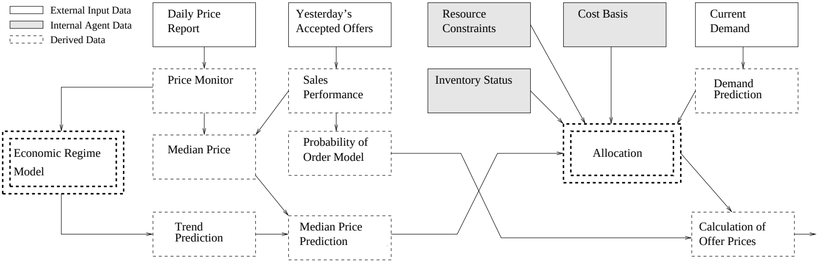

Related Work

Demand and Price Prediction

- Ghani, 2005 – PDA auctions on eBay
- Ghose et al., 2006 – used books sales on Amazon
- Kephart et al., 2000 – information goods and shopbots
- Massey et al., 2005 – reaction caused by regime shifts
- Osborn et al., 2002 – Macro-Economic regimes
- Pauwels et al., 2002 – windows of change in marketing

Demand and Price Prediction in TAC SCM

Benisch et al., 2004, Ketter et al., 2004, Pardoe et al., 2004, Wellman et al., 2005

Proposed Approach

1. Off-line
    1. Estimation of price density.
    2. Identification of regimes.
2. Real-time
    1. Identification of regimes.
    2. Prediction of regime distributions.
    3. Prediction of price density.
    4. Prediction of price trends.
    5. Prediction of order probability.
    6. Sales Pricing

Estimating Price Density Functions (1)

- Estimate price density functions and use them to define regimes.
- A Gaussian mixture model (GMM) can estimate arbitrary density functions.
- GMM is a semi-parametric approach:
  - fast computing
  - less memory

Estimating Price Density Functions (2)

We use a Gaussian mixture model (GMM):

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

where

- $p(\mathrm{np})$ is the density of the normalized price (np).
- $p(\mathrm{np}|\zeta_i) = N[\mu_i, \sigma_i](\mathrm{np})$ is the i-th Gaussian of the normalized price density from the GMM.
- $P(\zeta_i)$ is the prior probability of the i-th Gaussian.
- Fixed means μ_i and fixed variances σ_i^2.

Estimating Price Density Functions (3)

The EM-Algorithm determines the prior probability, $P(\zeta_i)$, of each Gaussian, where $\forall i = 1, \cdots, N$. Assumption: $N = 16$.

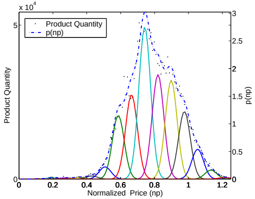

Low Market

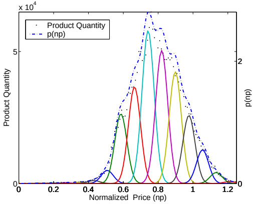

Medium Market

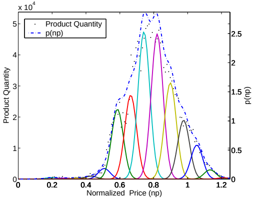

High Market

Using Bayes' rule we determine the posterior probability:

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

Proposed Approach

1. Off-line
   1. Estimation of price density.
   2. Identification of regimes.
2. Real-time
   1. Identification of regimes.
   2. Prediction of regime distributions.
   3. Prediction of price density.
   4. Prediction of price trends.
   5. Prediction of order probability.
   6. Sales Pricing

Definition of Regimes

We define the N-dimensional vector

$$\vec{\eta}(\mathrm{np}) = [P(\zeta_1|\mathrm{np}), P(\zeta_2|\mathrm{np}), \ldots, P(\zeta_N|\mathrm{np})]$$

1. Compute $\vec{\eta}(\mathrm{np}_j)$ which is $\vec{\eta}$ evaluated at the np_j price.
2. Cluster these collections of vectors using k-means.
3. The center of each cluster corresponds to a regime R_k.

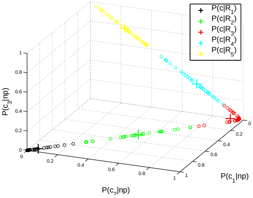

Low Market

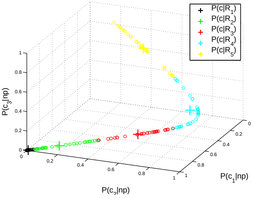

Medium Market

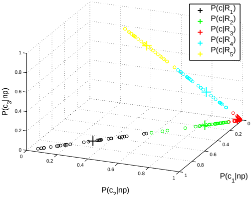

High Market

Off-line Regime Identification

Marginalizing over the components ζ_i we obtain:

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

where R_k is a specific regime.
Using Bayes' rule we determine the posterior probability:

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

The prior probabilities $P(R_k)$ are determined by a counting process over a collection of entire games.

Learned Regime Probabilities – offline

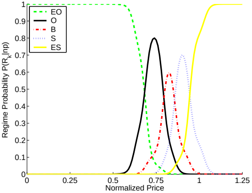

Low Market

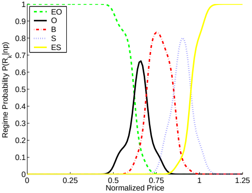

Medium Market

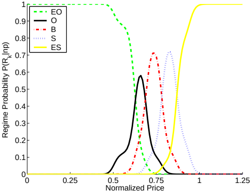

High Market

$\vec{P}(R_k|\mathrm{np}) \quad \forall k = 1, \cdots, M$ calculated off-line from 26 games.

Proposed Approach

1. Off-line
   1. Estimation of price density.
   2. Identification of regimes.
2. Real-time
   1. Identification of regimes.
   2. Prediction of regime distributions.
   3. Prediction of price density.
   4. Prediction of price trends.
   5. Prediction of order probability.
   6. Sales Pricing

Information Available in the Customer Market

Every day each agent receives:

1. Requests for Quotes (RFQs): computer type, number of computers, due date, reserve price.
2. A price report which includes the lowest and highest price paid per computer type from the previous day.

Online Identification of the Dominant Regime

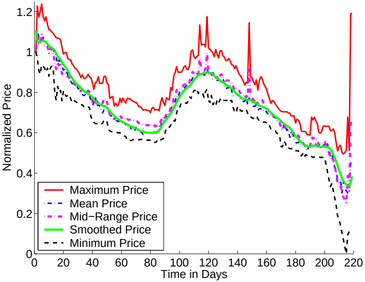

Daily price report 3721@tac3 – medium market: Minimum and maximum order prices.

1. Every day we estimate the current regime by calculating the double smoothed mid-range normalized price $\widetilde{\mathrm{np}}_{day}$ based on the daily price report.
2. We select the regime which has the highest probability, i.e. $\mathrm{argmax}_{1 \leq k \leq M} \vec{P}(R_k | \widetilde{\mathrm{np}}_{day})$.

Regime Probability – Real-time

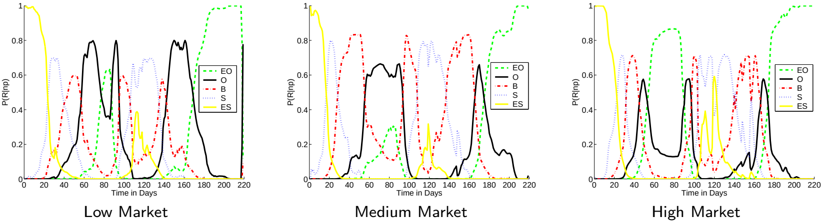

$\vec{P}(R_k|\widetilde{\mathrm{np}}_{day}) \quad \forall k = 1, \cdots, M$ calculated online for game 3721@tac3.

Regime Market Parameters (1)

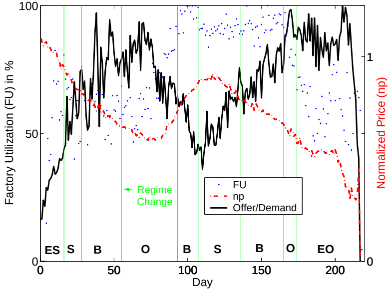

Factory Utilization (FU)

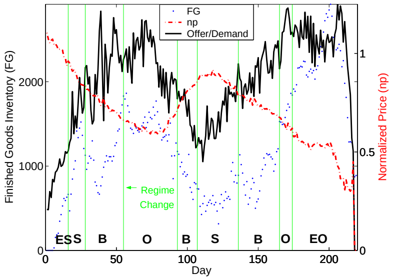

Finished Goods Inventory (FG)

Game 3721tac3 - medium market segment: Ratio offer/demand, normalized prices, and regime transitions.

Regime Market Parameters (2)

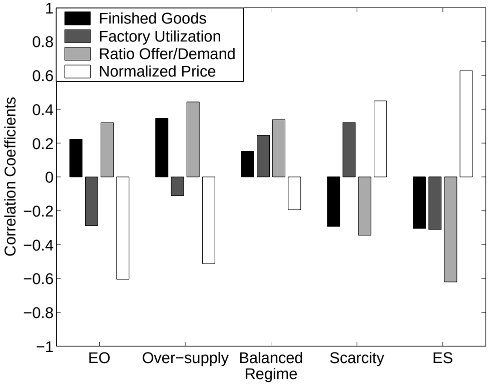

Training set (18 games) – Correlation coefficients between regimes and quantity of finished goods inventory, factory utilization, the ratio of offer to demand, and normalized price (np) in the medium market segment. All values are significant at the $p = 0.01$ level.

Proposed Approach

1. Off-line
   1. Estimation of price density.
   2. Identification of regimes.
2. Real-time
   1. Identification of regimes.
   2. Prediction of regime distributions.
   3. Prediction of price density.
   4. Prediction of price trends.
   5. Prediction of order probability.
   6. Sales Pricing

Online Prediction of Regimes (1)

We model the prediction of the next regime as a Markov prediction process: The posterior regime probabilities are predicted for current and future days based on yesterday's smoothed mid-range normalized price $\widetilde{\mathrm{np}}$.

Online Prediction of Regimes (2)

Repeated one-day prediction:

$$\vec{P}(r_{d+h}|\widetilde{\mathrm{np}}_{d-1}) = \sum_{r_{d+n}} \cdots \sum_{r_{d-1}} \left\{ \vec{P}(r_{d-1}|\widetilde{\mathrm{np}}_{d-1}) \cdot \mathbf{T}_{\mathbf{1}}^{\mathbf{h+1}}(r_d|r_{d-1}) \right\},$$

where

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

Proposed Approach

1. Off-line
    1. Estimation of price density.
    2. Identification of regimes.
2. Real-time
    1. Identification of regimes.
    2. Prediction of regime distributions.
    3. Prediction of price density.
    4. Prediction of price trends.
    5. Prediction of order probability.
    6. Sales Pricing

Prediction of Price Density

$$p(\widehat{\mathrm{np}}_{t+n}|\widetilde{\mathrm{np}}_{t-1}) = \sum_{j=1}^{N} P(\zeta_{j,t+n})\, p(\mathrm{np}|\zeta_j)$$

Sample np from 0 to 1.25 in increments of 0.01

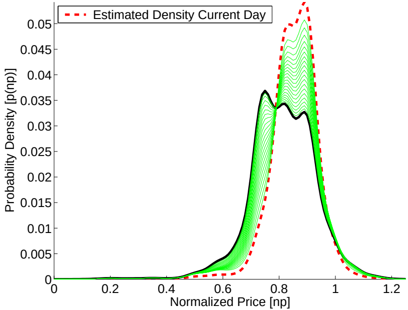

Example: Game 3717tac3

Evaluation of Price Prediction

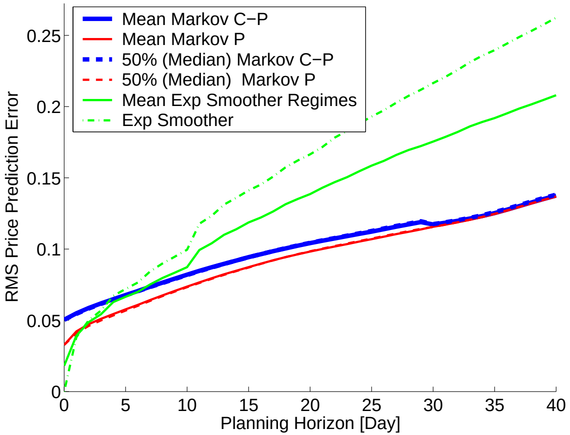

RMS error over a varying planning horizon.

Proposed Approach

1. Off-line
   1. Estimation of price density.
   2. Identification of regimes.
2. Real-time
   1. Identification of regimes.
   2. Prediction of regime distributions.
   3. Prediction of price density.
   4. Prediction of price trends.
   5. Prediction of order probability.
   6. Sales Pricing

Price Trend Prediction

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

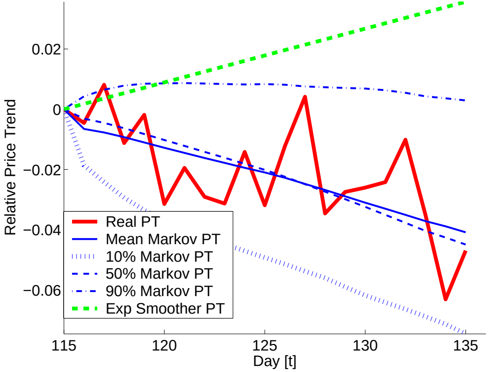

Example: Game 3717tac3

Evaluation of Trend Prediction

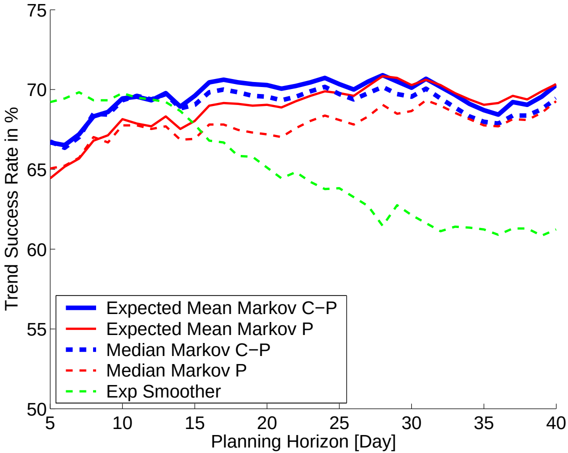

Success rate over a varying planning horizon.

Proposed Approach

1. Off-line
   1. Estimation of price density.
   2. Identification of regimes.
2. Real-time
   1. Identification of regimes.
   2. Prediction of regime distributions.
   3. Prediction of price density.
   4. Prediction of price trends.
   5. Prediction of order probability.
   6. Sales Pricing

Prediction of Order Probability

$$P(order|\mathrm{np}) = 1 - CDF(\mathrm{np})$$

Where the CDF is related to a probability density function $p(\mathrm{np})$ by

$$CDF(\mathrm{np}) = \int_0^{\mathrm{np}} p(\mathrm{np}')\,\mathrm{dnp}'$$

in the TAC SCM case $\mathrm{np}_{max} = 1.25$, so that $CDF(\mathrm{np}_{max}) = 1$.

Example Order Probability Curve

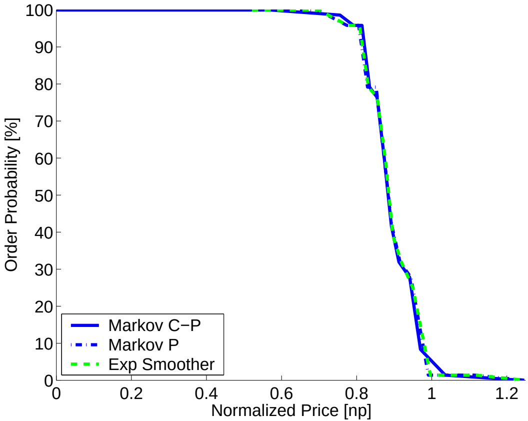

Real-time order probability curve for day 115 for the low market segment in game 3717@tac3.

Order Probability Prediction Results

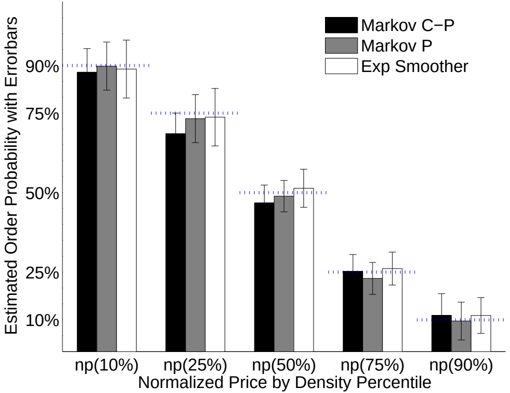

Daily order probability estimation (mean/std) for the 10th, 25th, 50th, 75th, and 90th percentile using different predictors.

Proposed Approach

1. Off-line
   1. Estimation of price density.
   2. Identification of regimes.
2. Real-time
   1. Identification of regimes.
   2. Prediction of regime distributions.
   3. Prediction of price density.
   4. Prediction of price trends.
   5. Prediction of order probability.
   6. Sales Pricing

Optimizing sales quotas (1)

To optimize profits over time, an agent needs to know:

- Current and future prices
- Its own costs
- Available inventory and production capacity

If per-unit profit for good g sold on day d at price $price_{d,g}$ is $\Phi_{d,g}$, then total profit over a horizon h is

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

Optimizing sales quotas (2)

LP solver can optimize total profit, subject to:

- Sales quotas cannot exceed expected demand
- Uncommitted finished-goods and raw-materials inventories
- Inventories are augmented by expected deliveries and components available from suppliers over the planning horizon h
- Quotas not satisfied from finished goods are constrained by factory capacity

Setting sales prices

To sell our quota, we set prices using our order probability model.

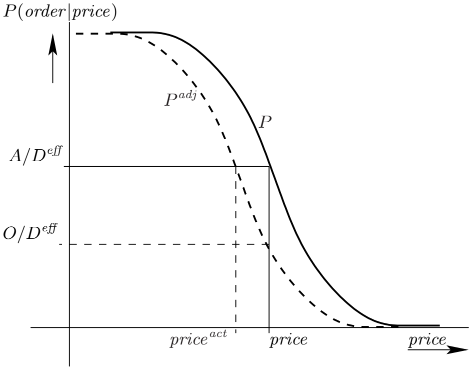

Tuning sales prices

The pricing model is approximate. We tune it using feedback from actual orders.

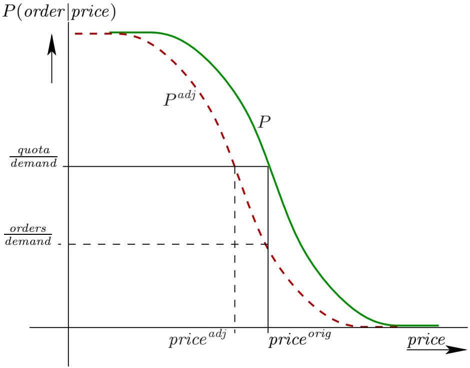

Mean Profit Results using Different Versions of MinneTAC

	Mean Profit / Standard deviation (in $M)			
Strategic:	Price-Follower	Regimes	Combo	Regimes
Tactical:	Linear	Linear	Linear	Regimes
Agent:				
MinneTAC	1.347/3.703	1.813/4.017	1.780/4.536	2.117/3.764
TacTex06	8.752/5.682	8.873/5.600	8.399/5.173	9.205/5.385
DeepMaize06F	8.839/4.629	8.713/4.846	8.403/4.710	8.318/4.181
PhantAgent06	8.049/5.422	7.991/5.384	7.895/5.326	8.173/5.437
Maxon06F	4.243/4.516	3.767/4.288	3.808/4.254	4.019/4.181
Rational05	0.739/4.912	0.669/4.692	0.710/4.692	1.305/4.527

Experimental setup with controlled market conditions and different variations of MinneTAC for order probability, price and price trend predictions. Each column is an average of 23 games.

Future Work (1)

- Ensemble price predictions
- Train regime transition matrices:
  - On different time periods (start, mid, and end of the game).
  - Include the effect of substitutability among market segments and products.
- Market segments vs product learning.
- Develop procurement strategies that take advantage of regime forecasting.

Future Work (2)

- Integrate regime forecasting in decision making process. Apply reinforcement learning to map
  - economic regimes to operational regimes.
  - operational regimes to actions.
- Implement and evaluate approach in other application domains, e.g.,
  - Stock market
  - Amazon
  - Dutch flower auction
  - Energy markets

Conclusions

- Off-line identification of economic regimes.
- Real-time identification of economic regimes.
- Real-time prediction of economic regime distributions and transitions.
- Real-time prediction of price density, prices and price trends.
- Real-time prediction of order probability.

Contact

email: `wketter@rsm.nl`

URL: `www.ketter.ws`

Prediction of Price Density

$$
\begin{aligned}
&p(\widehat{\mathrm{np}}_{t+n}|\widetilde{\mathrm{np}}_{t-1}) \\
&= \sum_{i=1}^{M} P(\mathrm{np}|R_i)\, P(R_{i,t+n}|\widetilde{\mathrm{np}}_{t-1}) \\
&= \sum_{j=1}^{N} \underbrace{\sum_{i=1}^{M} P(\zeta_j|R_i)\, P(R_{i,t+n}|\widetilde{\mathrm{np}}_{t-1})}_{P(\zeta_{j,t+n})}\, p(\mathrm{np}|\zeta_j) \\
&= \sum_{j=1}^{N} P(\zeta_{j,t+n})\, p(\mathrm{np}|\zeta_j)
\end{aligned}
$$