Identification and Prediction of Economic Regimes to Guide Decision Making in Multi-Agent Marketplaces

Wolfgang Ketter

Department of Computer Science and Engineering

University of Minnesota

Adventium Labs

Minneapolis - January 17th 2007

Overview

- Motivation
- Trading Agent Competition for Supply Chain Management (TAC SCM)
- Related Work
- Proposed Solution
- Future Work
- Conclusion

Motivation

- Economic theory suggests that economic environments exhibit 3 dominant market patterns: scarcity, balanced, and over-supply.
- We call these distinguishable conditions *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 between Prices, Order Probability, and Regimes

Scarcity:

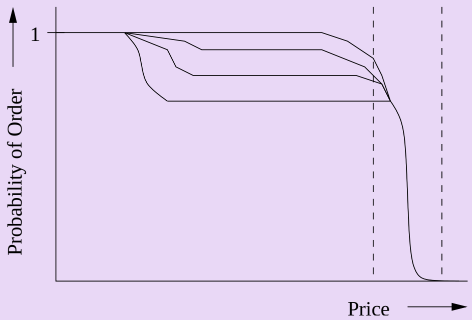

Balanced:

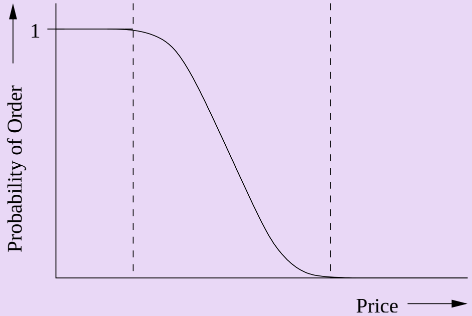

Over-supply:

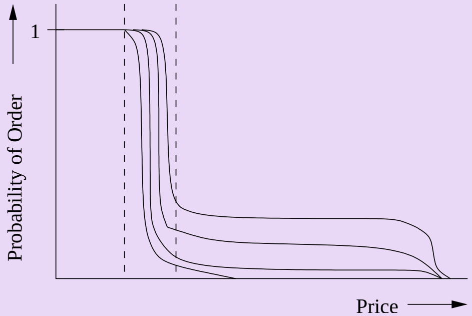

Reverse cumulative density function represents probability of order.

Experimental:

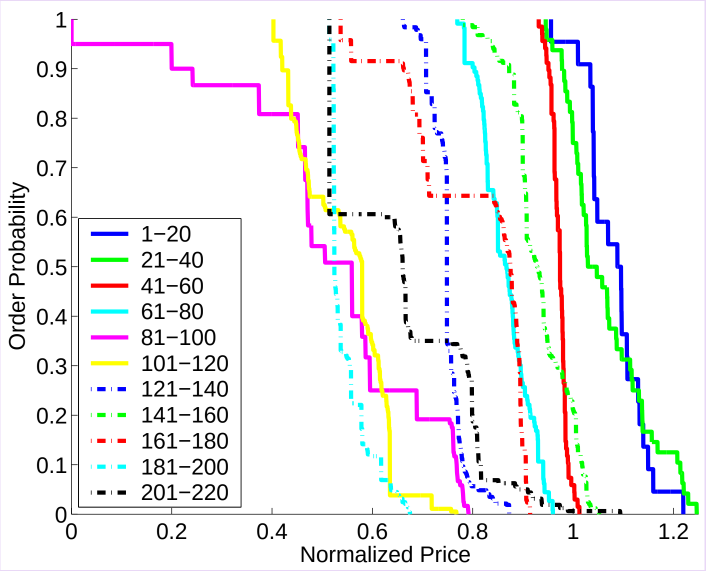

Application Areas (1)

- 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-Chrysler).
- Automated supply-chain management, e.g.,
  - i2
  - SAP

Application Areas (2)

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

TAC SCM - Scenario

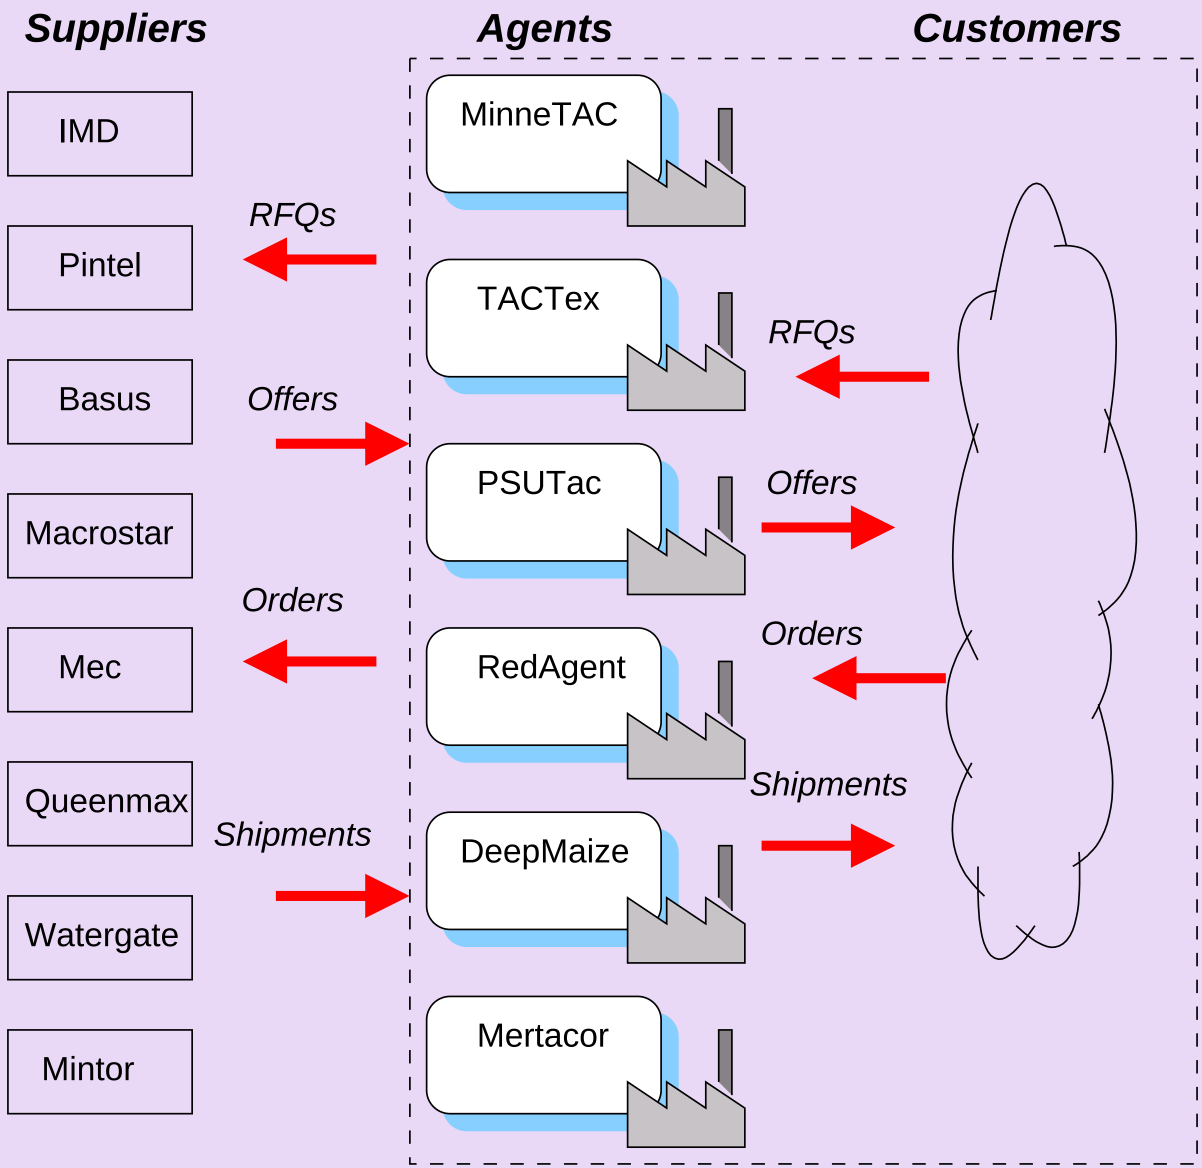

Use Regime Prediction For Sales Strategies

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

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
   (a) Estimation of price density.
   (b) Identification of regimes.
2. Real-time
   (a) Identification of regimes.
   (b) Prediction of regime distributions.
   (c) Prediction of price density.
   (d) Prediction of price trends.
   (e) Prediction of order probability.

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: $N = 16$.

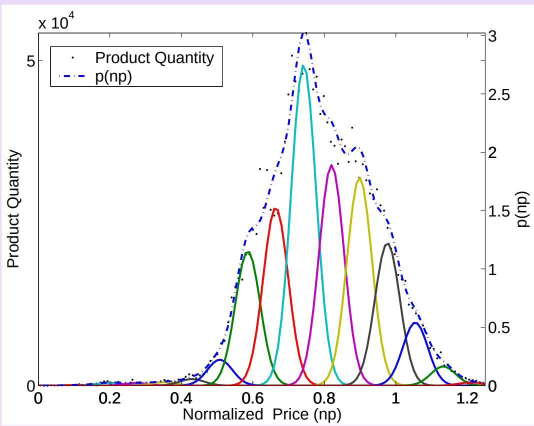

Low Market

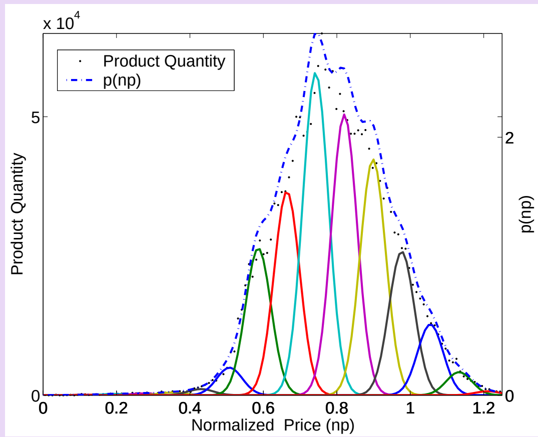

Medium Market

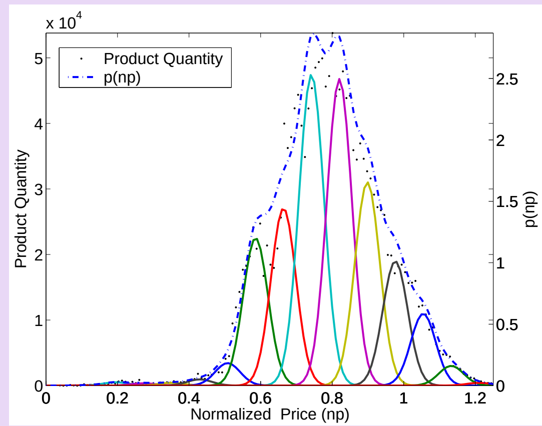

High Market

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
   (a) Estimation of price density.
   (b) Identification of regimes.
2. Real-time
   (a) Identification of regimes.
   (b) Prediction of regime distributions.
   (c) Prediction of price density.
   (d) Prediction of price trends.
   (e) Prediction of order probability.

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.

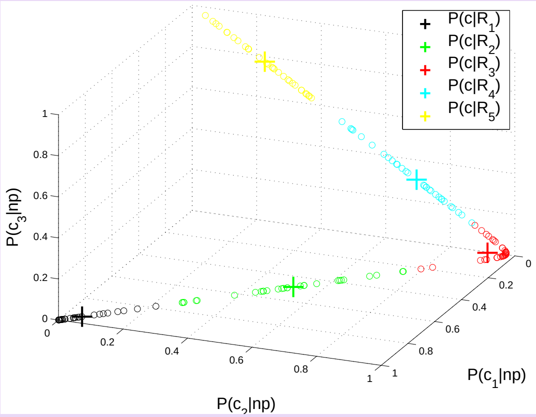

Low Market

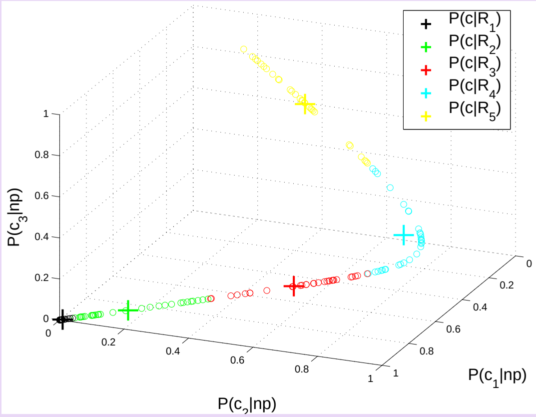

Medium Market

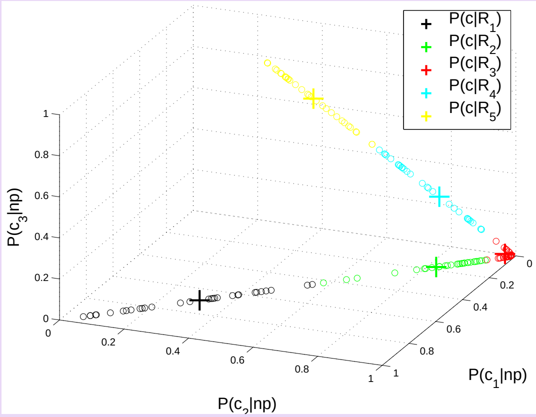

High Market

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

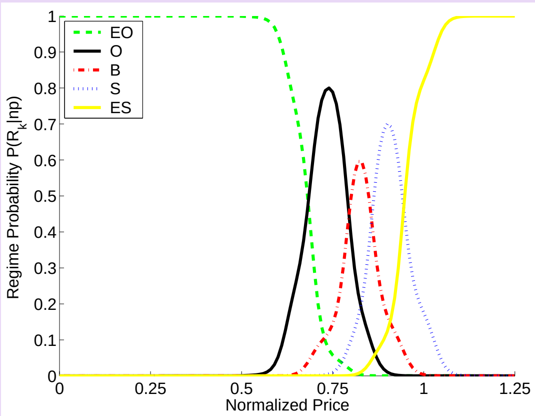

Low Market

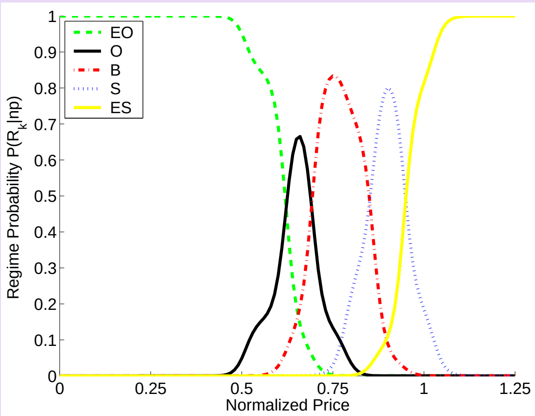

Medium Market

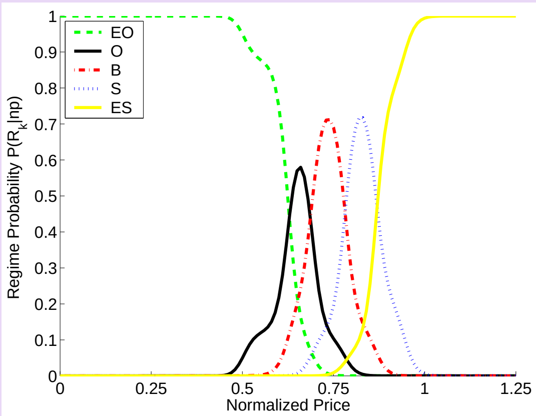

High Market

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

Proposed Approach

1. Off-line
   (a) Estimation of price density.
   (b) Identification of regimes.
2. Real-time
   (a) Identification of regimes.
   (b) Prediction of regime distributions.
   (c) Prediction of price density.
   (d) Prediction of price trends.
   (e) Prediction of order probability.

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 Current Regime

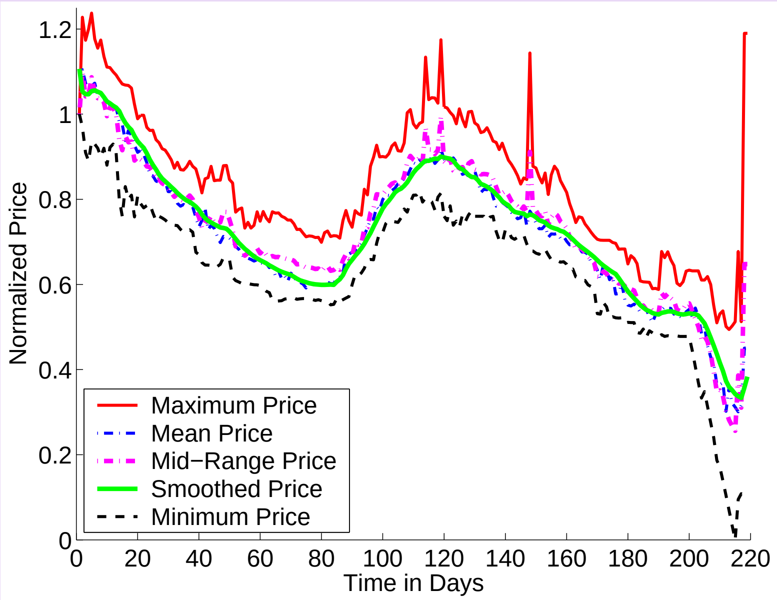

Daily price report 3721@tac3 – medium market: Minimum and maximum order prices.

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. $\mathrm{argmax}_{1 \le k \le M} \vec{P}(R_k | \widetilde{\mathrm{np}}_{day})$.

Example Regime Probability – Real-time

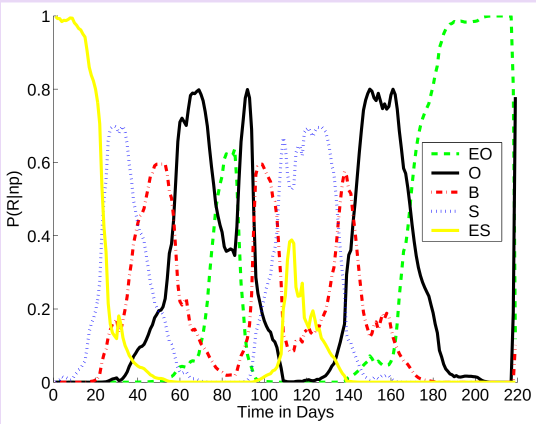

Low Market

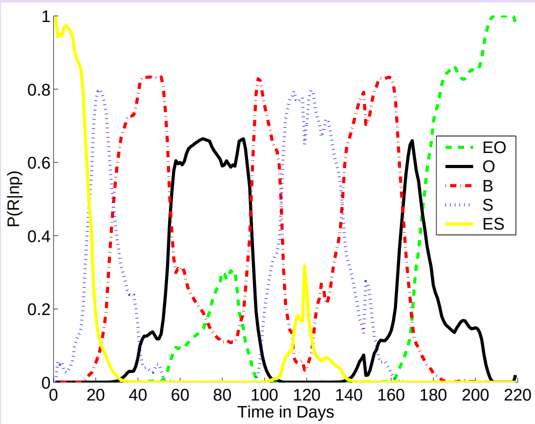

Medium Market

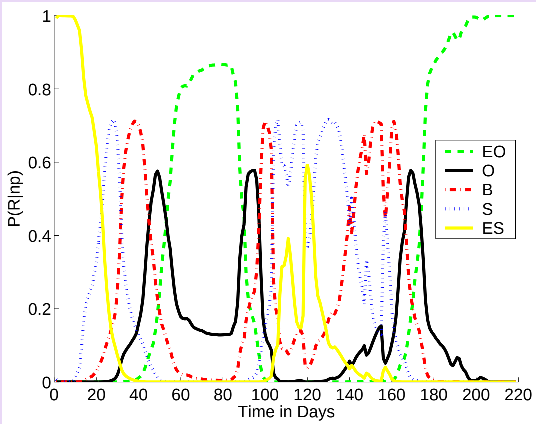

High Market

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

Regime Market Parameters (1)

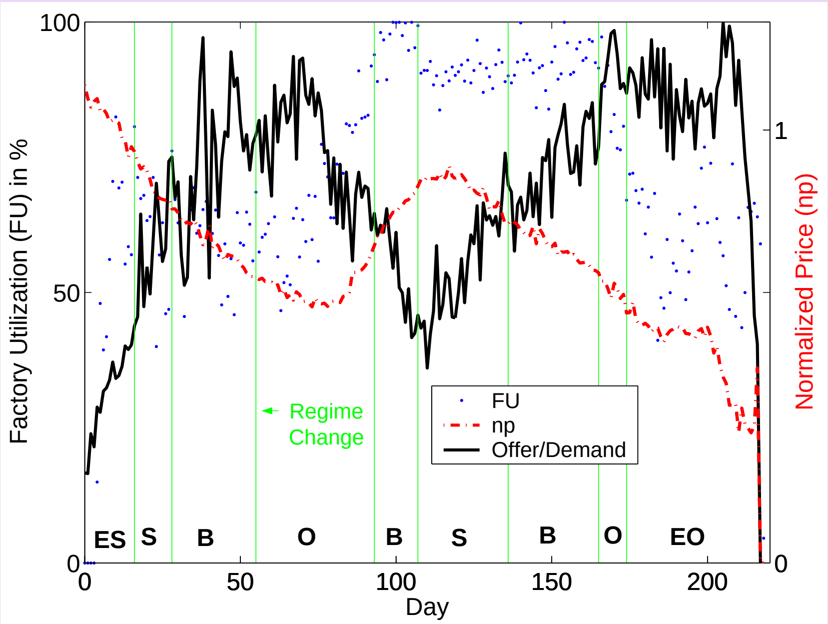

Factory Utilization (FU)

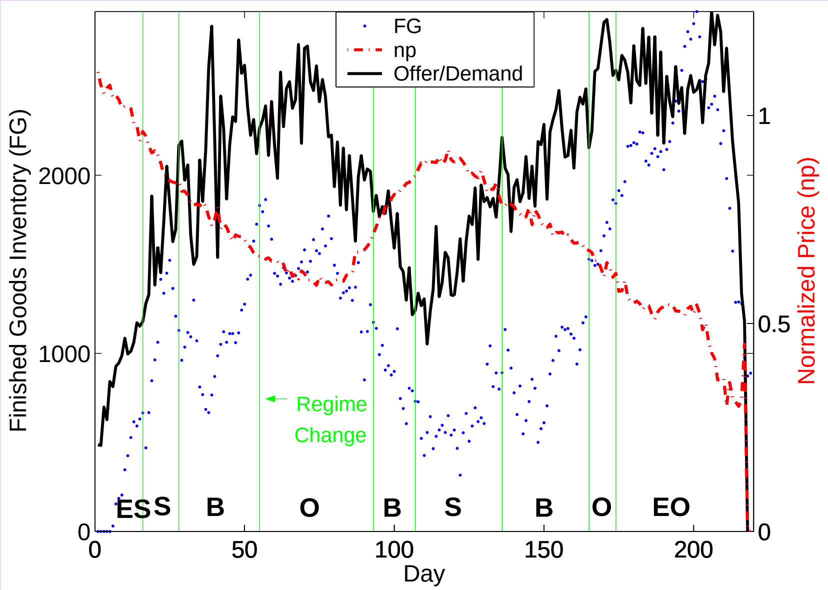

Finished Goods Inventory (FG)

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

Regime Market Parameters (2)

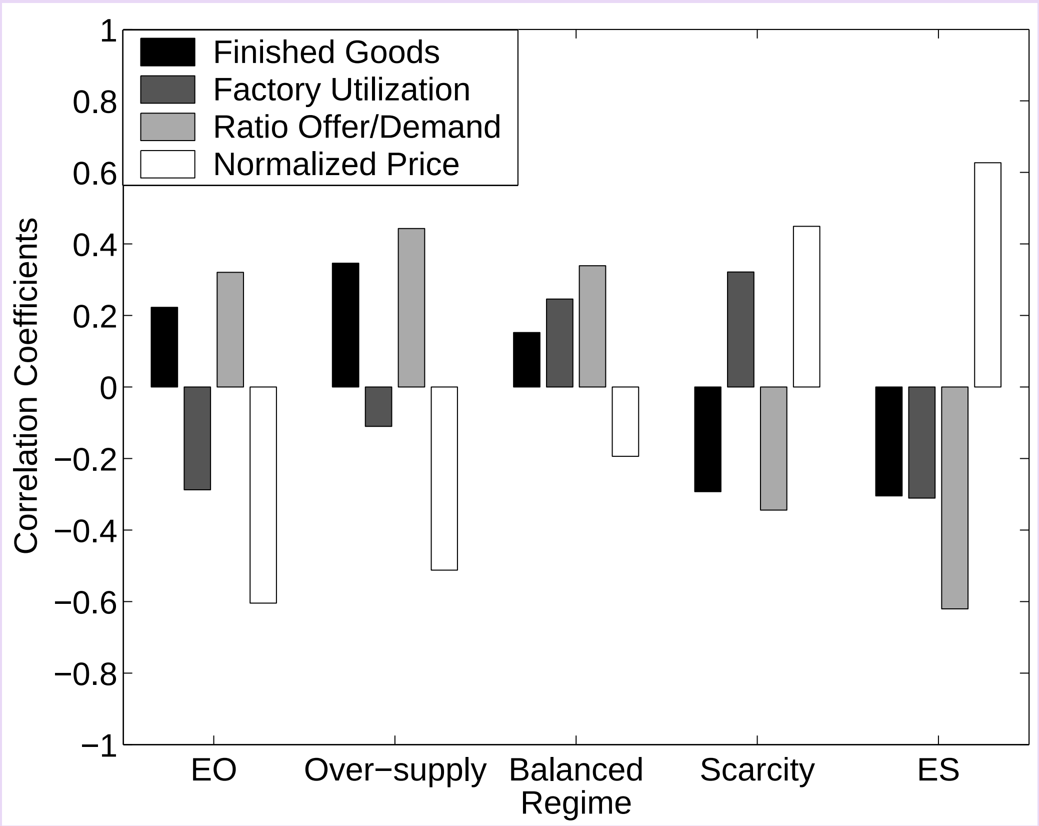

Training set (18 games) – Correlation coefficients between regimes and quantity of finished goods inventory, factory utilization, the ratio of offer to demand, and normalized price (np) in the medium market segment. All values are significant at the $p = 0.01$ level.

Proposed Approach

1. Off-line
   (a) Estimation of price density.
   (b) Identification of regimes.
2. Real-time
   (a) Identification of regimes.
   (b) Prediction of regime distributions.
   (c) Prediction of price density.
   (d) Prediction of price trends.
   (e) Prediction of order probability.

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:

$$\vec{P}(r_{d+h}|\widetilde{\mathrm{np}}_{d-1})$$
$$= \sum_{r_{d+n}} \cdots \sum_{r_{d-1}} \left\{ \vec{P}(r_{d-1}|\widetilde{\mathrm{np}}_{d-1}) \cdot \mathbf{T_1^{h+1}}(r_d|r_{d-1}) \right\},$$

where

$$\mathbf{T_1^{h+1}}(r_d|r_{d-1}) = \prod_{n=0}^{h} \mathbf{T_1}(r_d|r_{d-1})$$

Prediction of Dominant Regime

Prediction results computed every day for the next 40 days from day 1 to day 179 for 5 regimes.

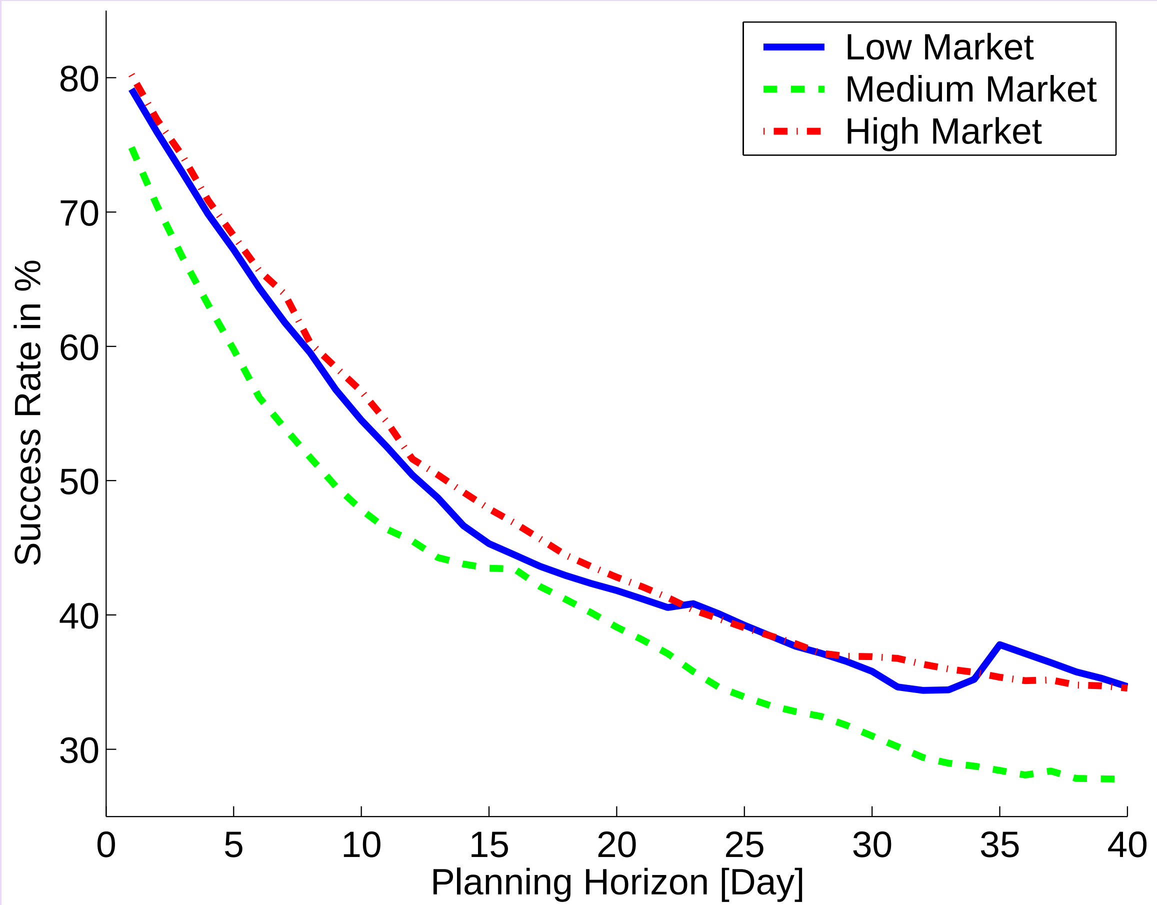

Compare Regime Distributions

Kullback-Leibler (KL) divergence:

- measures the difference between two probability distributions in bits
- the smaller the measure the closer the predictions are to optimal

$$KL(\vec{P}(\hat{R})\|\vec{P}(R)) = \sum_{r \in \mathcal{R}} \vec{P}(\hat{R}) \log \left(\frac{\vec{P}(\hat{R})}{\vec{P}(R)} \right)$$

where $\vec{P}(\hat{R})$ is the predicted regime distribution and $\vec{P}(R)$ is the actual regime distribution.

Prediction of Regime Distribution

Prediction results computed every day for the next 40 days from day 1 to day 179 for 5 regimes.

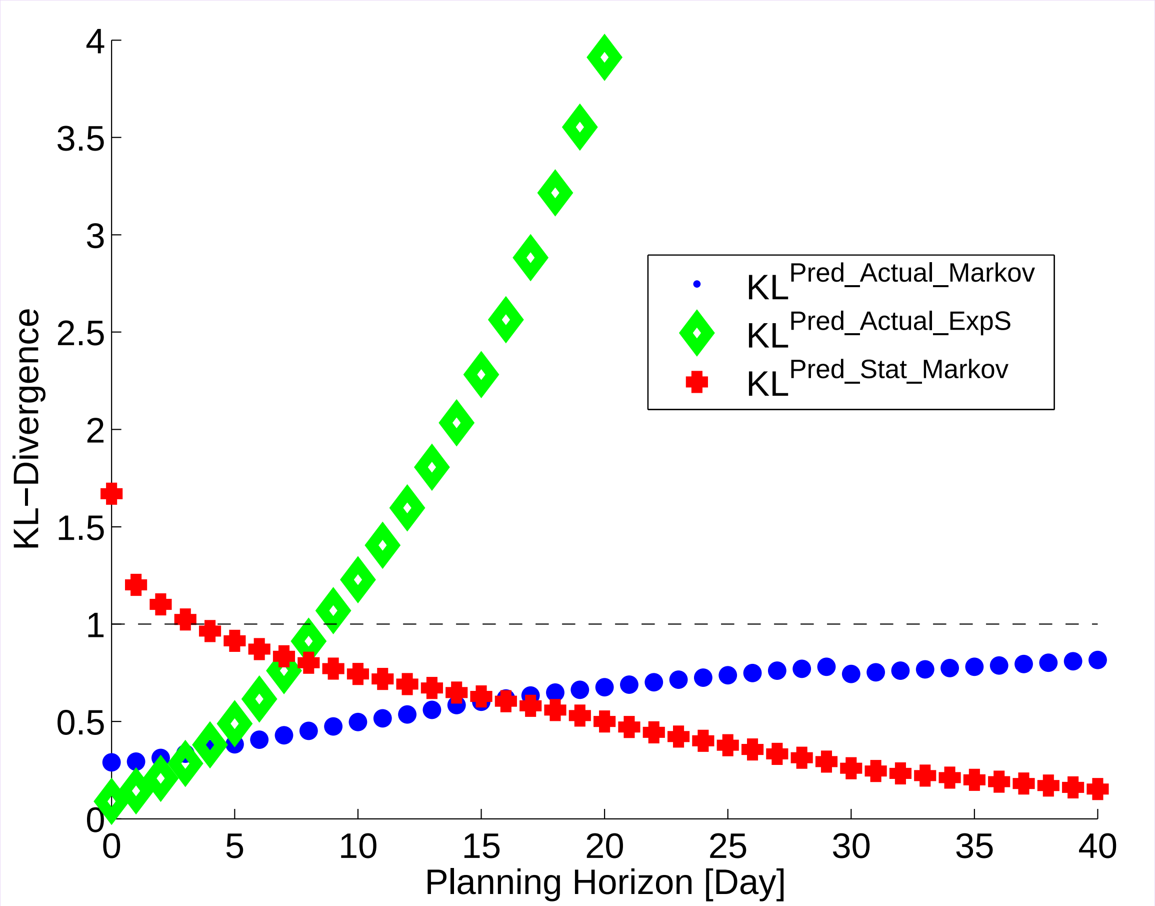

Proposed Approach

1. Off-line
   (a) Estimation of price density.
   (b) Identification of regimes.
2. Real-time
   (a) Identification of regimes.
   (b) Prediction of regime distributions.
   (c) Prediction of price density.
   (d) Prediction of price trends.
   (e) Prediction of order probability.

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

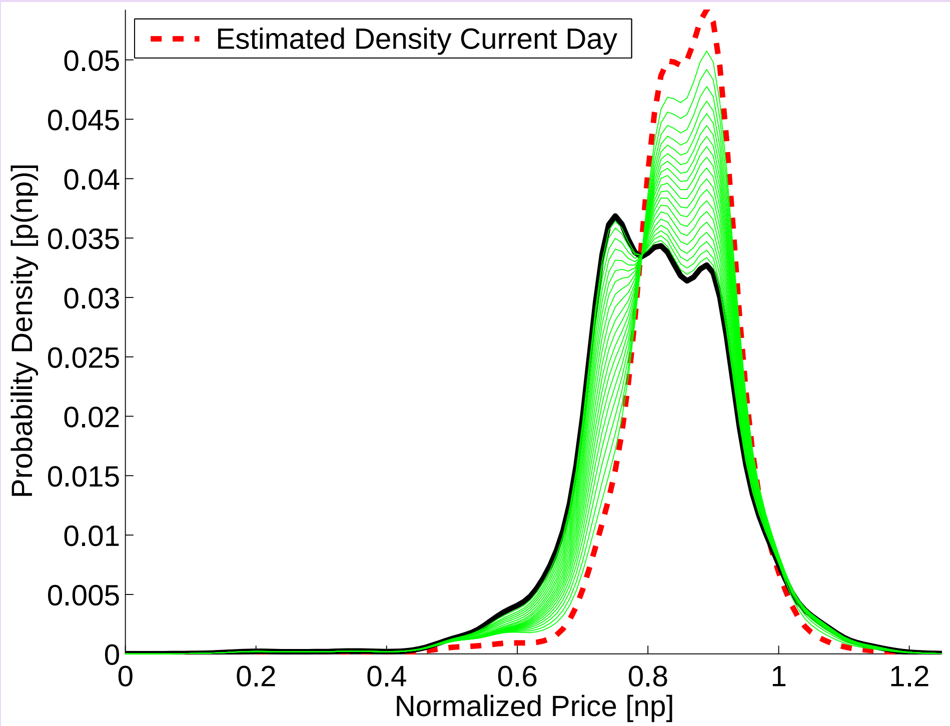

$\Rightarrow$

Evaluation of Price Prediction

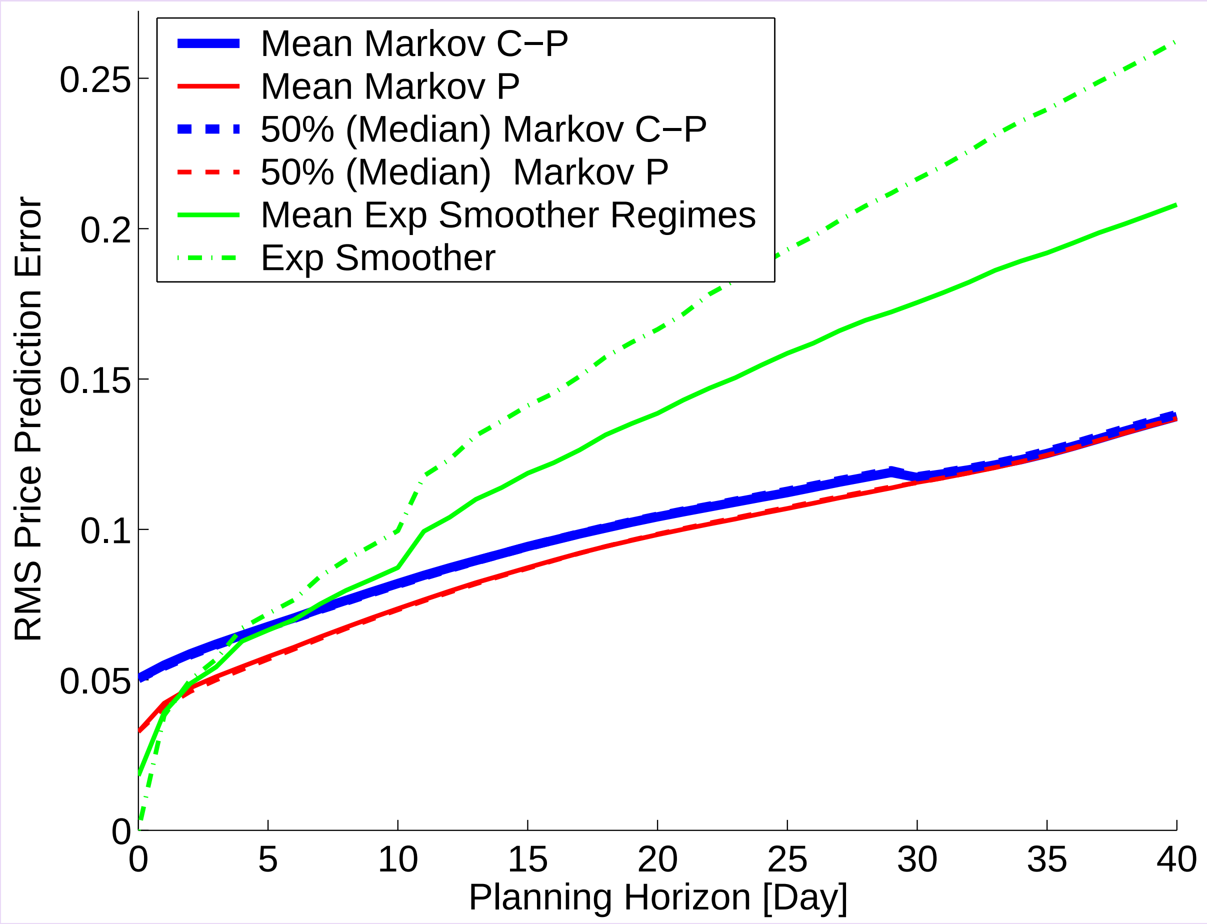

RMS error over a varying planning horizon.

Proposed Approach

1. Off-line
   (a) Estimation of price density.
   (b) Identification of regimes.
2. Real-time
   (a) Identification of regimes.
   (b) Prediction of regime distributions.
   (c) Prediction of price density.
   (d) Prediction of price trends.
   (e) Prediction of order probability.

Price Trend Prediction

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

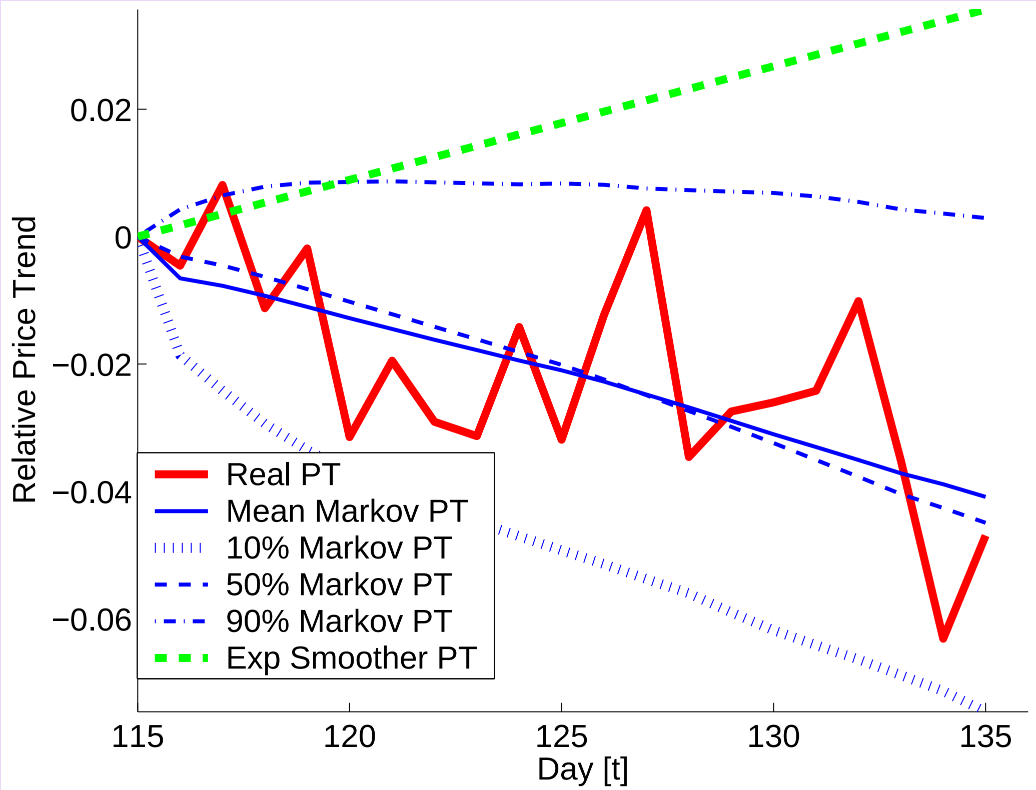

Example: Game 3717tac3

Evaluation of Trend Prediction

Success rate over a varying planning horizon.

Proposed Approach

1. Off-line
   (a) Estimation of price density.
   (b) Identification of regimes.
2. Real-time
   (a) Identification of regimes.
   (b) Prediction of regime distributions.
   (c) Prediction of price density.
   (d) Prediction of price trends.
   (e) Prediction of order probability.

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

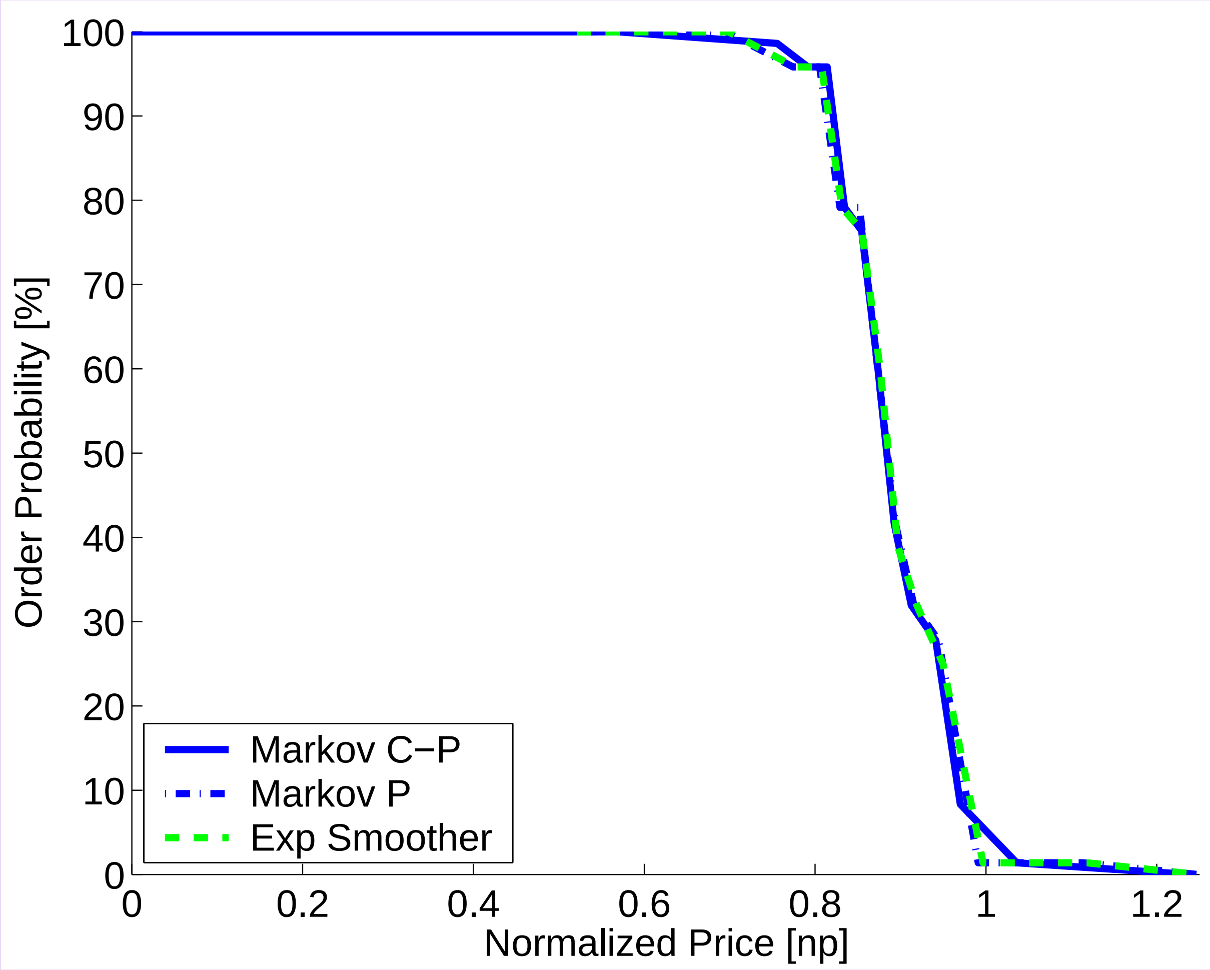

Real-time order probability curve for day 115 for the low market segment in game 3717@tac3.

Order Probability Prediction Results

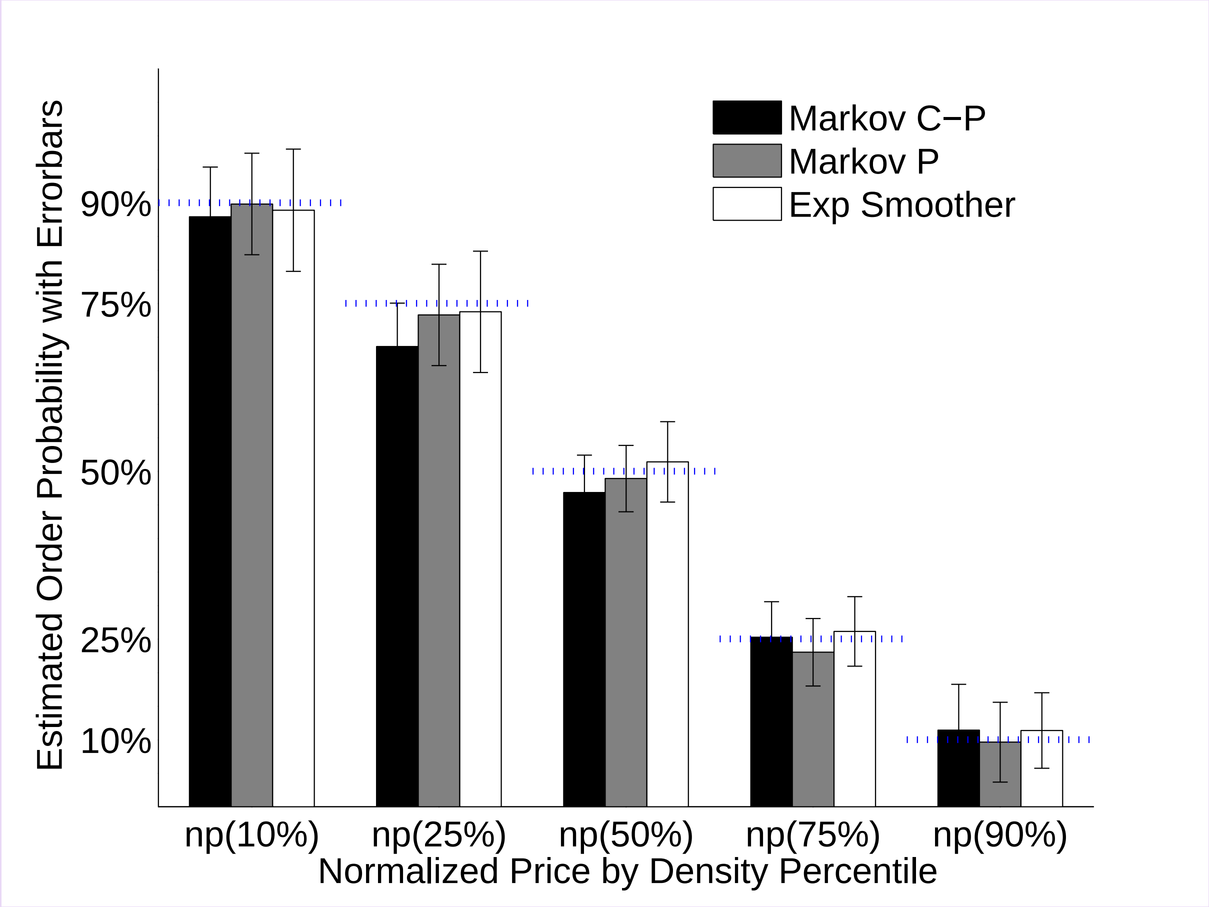

Daily order probability estimation (mean/std) for the 10th, 25th, 50th, 75th, and 90th percentile using different predictors.

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)

- Enhance sales strategy to take full advantage of regime forecasting and price density, e.g.,
  - quantify regime prediction
  - measure confidence interval
- Develop procurement strategies that take advantage of regime forecasting.
- 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.

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.,
  - Amazon
  - eBay

Conclusions

- Off-line identification of economic regimes from past game data.
- Online identification of economic regimes from data available in the current game.
- Prediction of economic regime transitions.
- Prediction of price density and price trends.
- Prediction of order probability.

Contact
email: ketter@cs.umn.edu
URL: www.ketter.ws

Credits

Many thanks to:

- John Collins
- Maria Gini
- Alok Gupta
- Paul Schrater

Prediction of Price Density

$$p(\widehat{\mathrm{np}}_{t+n} | \{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{t-1} \ )$$

$$= \sum_{i=1}^{M} P(\mathrm{np}|R_i)\, P(R_{i,t+n}|\{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{t-1} \ )$$

$$= \sum_{j=1}^{N} \underbrace{\sum_{i=1}^{M} P(\zeta_j|R_i)\, P(R_{i,t+n}|\{\widetilde{\mathrm{np}}_1, \ldots, \widetilde{\mathrm{np}}_{t-1} \ )}_{P(\zeta_{j,t+n})} p(\mathrm{np}|\zeta_j)$$

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

$\Lleftarrow$