Market Making Recap

- A market maker is always willing to trade
- Sets bid and ask prices: $b_t < a_t$
- May have to take risks, hold inventory, trade against more informed people, etc.
- De facto standard in prediction markets: the Logarithmic Market Scoring Rule (LMSR) (Hanson 2003, 2007; Chen & Pennock 2007)
  - Loss-making, but bounded loss
  - Many nice theoretical properties
  - Can be unstable, depending on a key parameter

Bid and Ask Prices

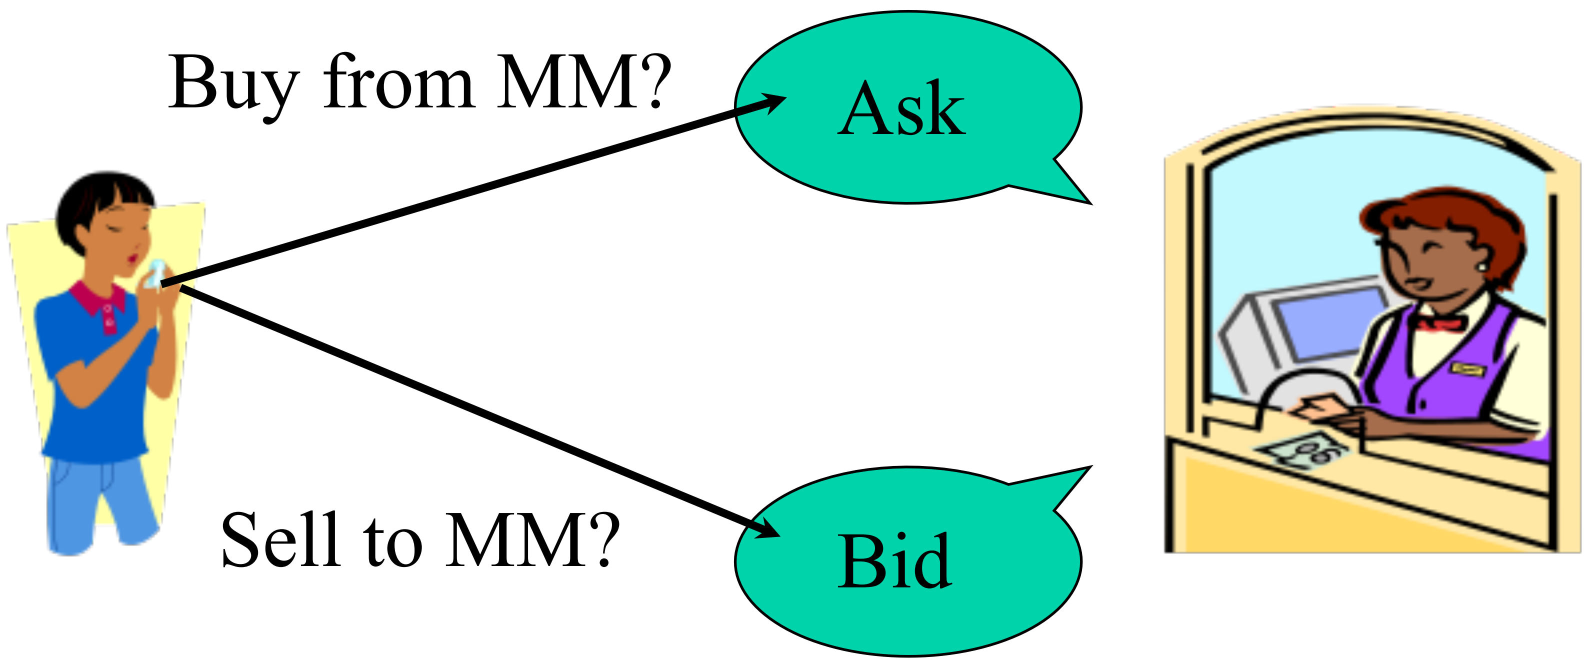

- Market-maker faces a censored learning problem
- Two parts
  - Inference given bid and ask prices
  - How to set bid and ask prices

A Reinforcement Learning Perspective

Prices are actions; trades are signals

A specific model: universe of traders with normally distributed beliefs about the true value (arrive sequentially)

$$w_t = V + \epsilon_t$$

$$\text{Buy if } w_t > a_t, \text{ sell if } w_t < b_t$$

MM attempts to infer V
-- State space is functional: belief distribution on V

[Das, Quant. Fin 05, AAMAS 08]

Inference: State Space Updates

Let F_ϵ be the c.d.f. of the valuation distribution.

Then, outcome probabilities are:

$$q_t(V; b_t, a_t) = \begin{cases} 1 - F_\epsilon(a_t - V) \\ F_\epsilon(a_t - V) - F_\epsilon(b_t - V) \\ F_\epsilon(b_t - V) \end{cases}$$

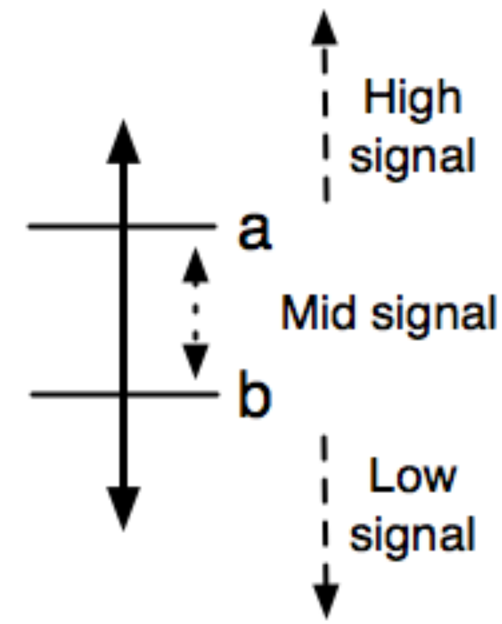

Bayesian update

$$p_{t+1}(v) = p_t(v) \frac{q_t(v; b_t, a_t)}{\mathcal{A}_t}$$

Inference: Computing the Next State

Gaussian prior
Problem: not conjugate for censored observations
-- Posterior is no longer Gaussian

Enforce Gaussian-ness: extract mean and variance of the true updated distribution and match moments

$$\mu_{t+1} = \mu_t + \sigma_t \cdot \frac{B}{A}$$

$$\sigma_{t+1}^2 = \sigma_t^2 \left(1 - \frac{AC + B^2}{A^2}\right)$$

(A, B, and C are integrals of various forms of Gaussians)

State update is monotonic!

$$\sigma_{t+1}^2 \leq \sigma_t^2$$

[Das & Magdon-Ismail, NIPS 08]

Price Setting

Profit maximizing strategies

-- Monopolist MM (NYSE) => long-term (sequential) profit maximization

-- Competitive MMs (NASDAQ) => zero-expected profit pricing (Glosten & Milgrom, 1985; Kyle, 1985)

+ market designer could design it this way for liquidity provision

$$b_t = \mathbb{E}[V|\text{Sell}] = \frac{\int dv \quad v p_t(v) F_\epsilon(b_t - v)}{\int dv \quad p_t(v) F_\epsilon(b_t - v)}$$

In the Gaussian case:

$$b_t = \mu_t - q\sigma_\epsilon \sqrt{1 + \rho^2}$$

where $\rho = \dfrac{\sigma}{\sigma_\epsilon}$ and $q = \dfrac{\rho^2}{1+\rho^2} \dfrac{N(q)}{1 - \Phi(q)}$

Price Setting: Optimal Monopolist

- Exploration vs. exploitation: prices as sources of both information and profit
- Monotonic variance update: allows efficient, single-sweep DP solution

σ	0	0.01	0.02	0.03	...	9.98	9.99	10.0
V	V(0)							

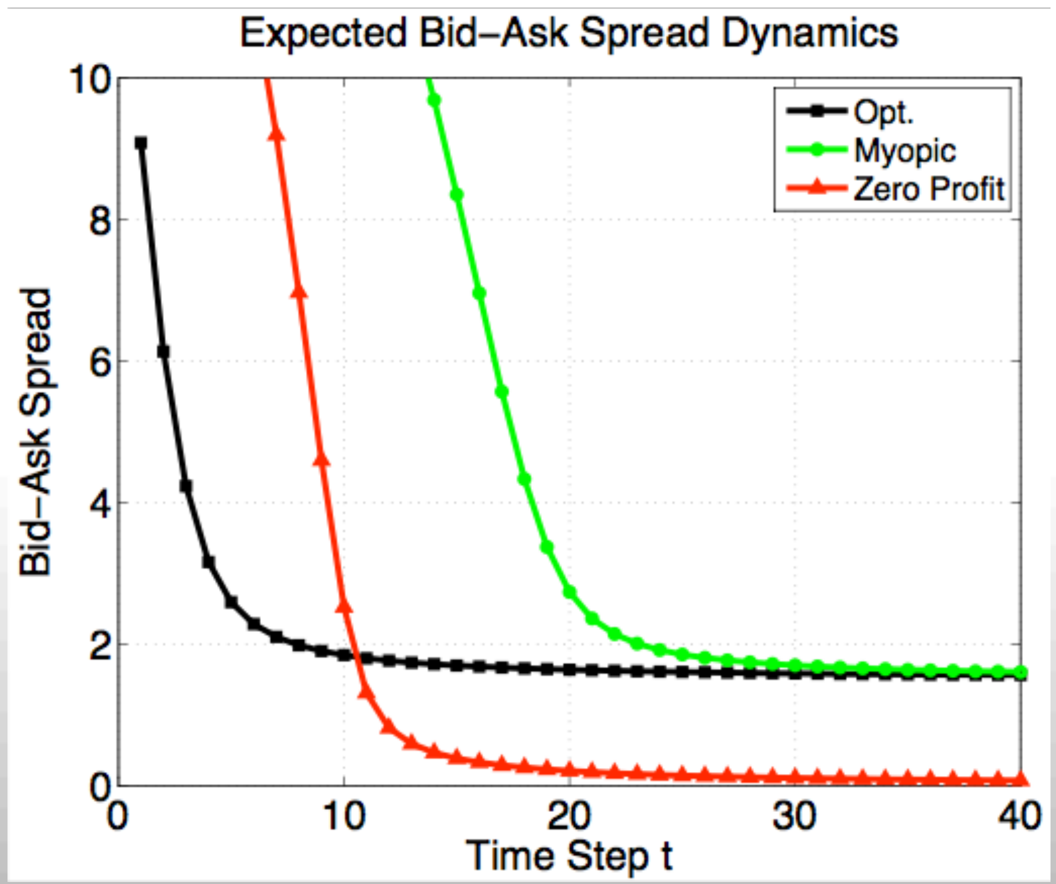

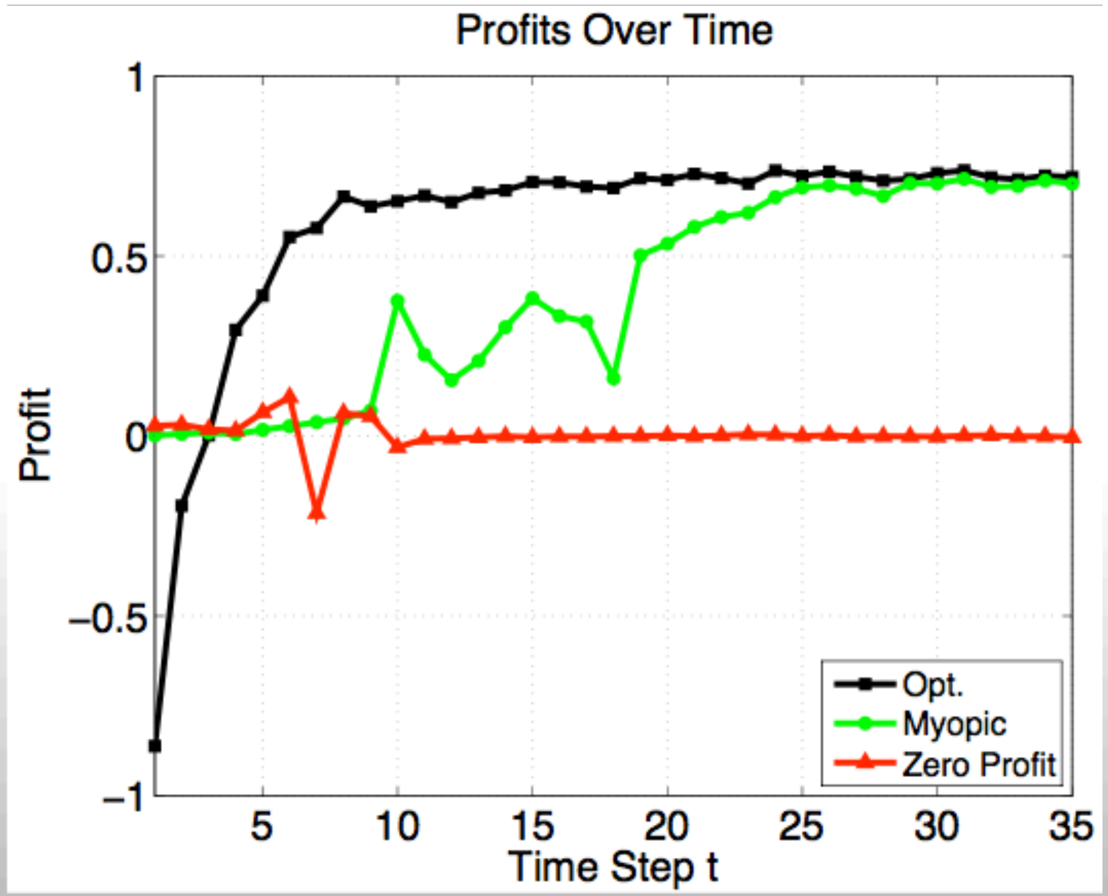

Discussion

- NYSE and NASDAQ used to debate the merits of monopoly vs competition between market makers.
  - NYSE: a monopolist can ``maintain a fair and orderly market" in the face of market shocks

- Glosten (1989) showed that monopolists can provide greater liquidity under asymmetric information by averaging expected profits across different trade sizes.

- We show that this can hold true with fixed trade sizes in a multi-period setting because of the exploration/ exploitation tradeoff

3 Experiments

- For validation in practice: compare with the standard LMSR MM in various scenarios
  - 1. Human subjects playing a 10-15 minute trading game ("Bouncing Balls")
  - 2. Trading agents competing in a simulated market ("Trading Agents")
  - 3. Human traders participating in a months-long field experiment ("Instructor Ratings")

A Short Detour

Making the market maker practical involves making it adaptive to jumps in the true value

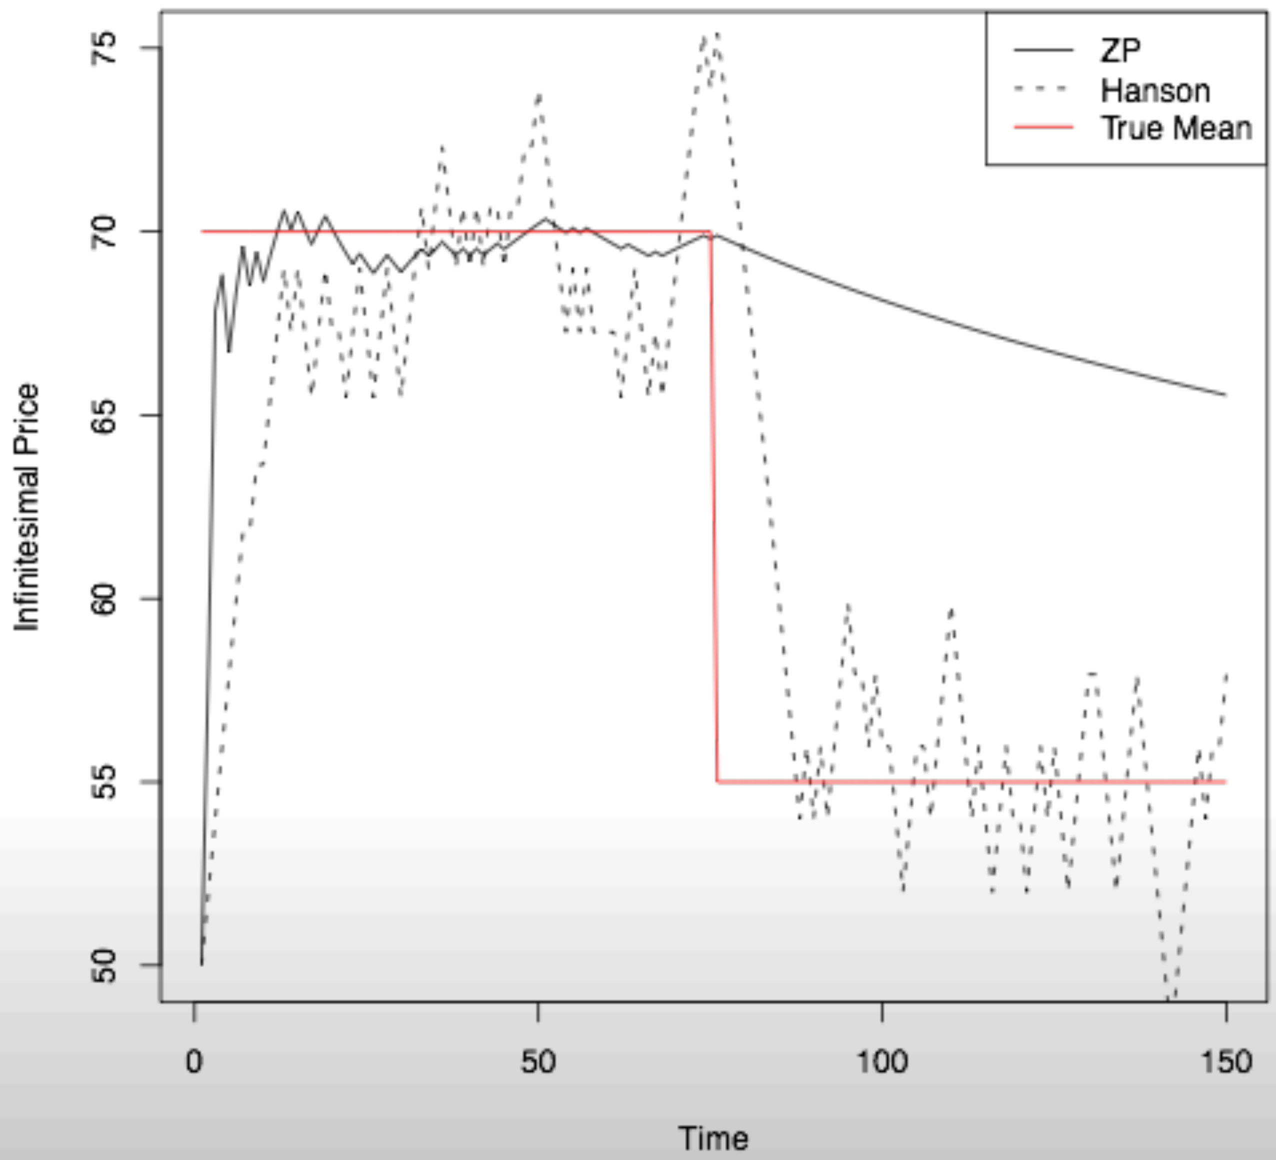

A Solution

Prob. of an observed sequence of trades

$$L(\mu, \sigma) = \int_{-\infty}^{\infty} N(v, \mu, \sigma) . \prod_{i=1}^{s} \left(\Phi(z_i^+, v, \sigma_\epsilon) - \Phi(z_i^-, v, \sigma_\epsilon) \right) dv$$

Was this sequence more likely under MM's current belief variance or an increased variance?

- If the latter, double the variance of current belief

Big benefits:

- Relative measure, not too sensitive to window size
- Rapid adaptivity

[Brahma, Chakraborty, Das, Lavoie & Magdon-Ismail, ACM EC 2012]

Evaluation Metrics

- Profit/Loss
  - Average
  - Maximum loss (or other risk measures)

- Liquidity provided (measured by spread)

- "Correctness" of the markets
  - RMSD of prices versus "true value"
  - RMSD in equilibrium phase of the market (after price convergence)

Experiment 1: Bouncing Balls

- Humans trade simultaneously in two markets, without knowing which market maker is involved in each one
- They trade on the proportion of times a ball will fall off one or the other edge in a "gambler's ruin" game

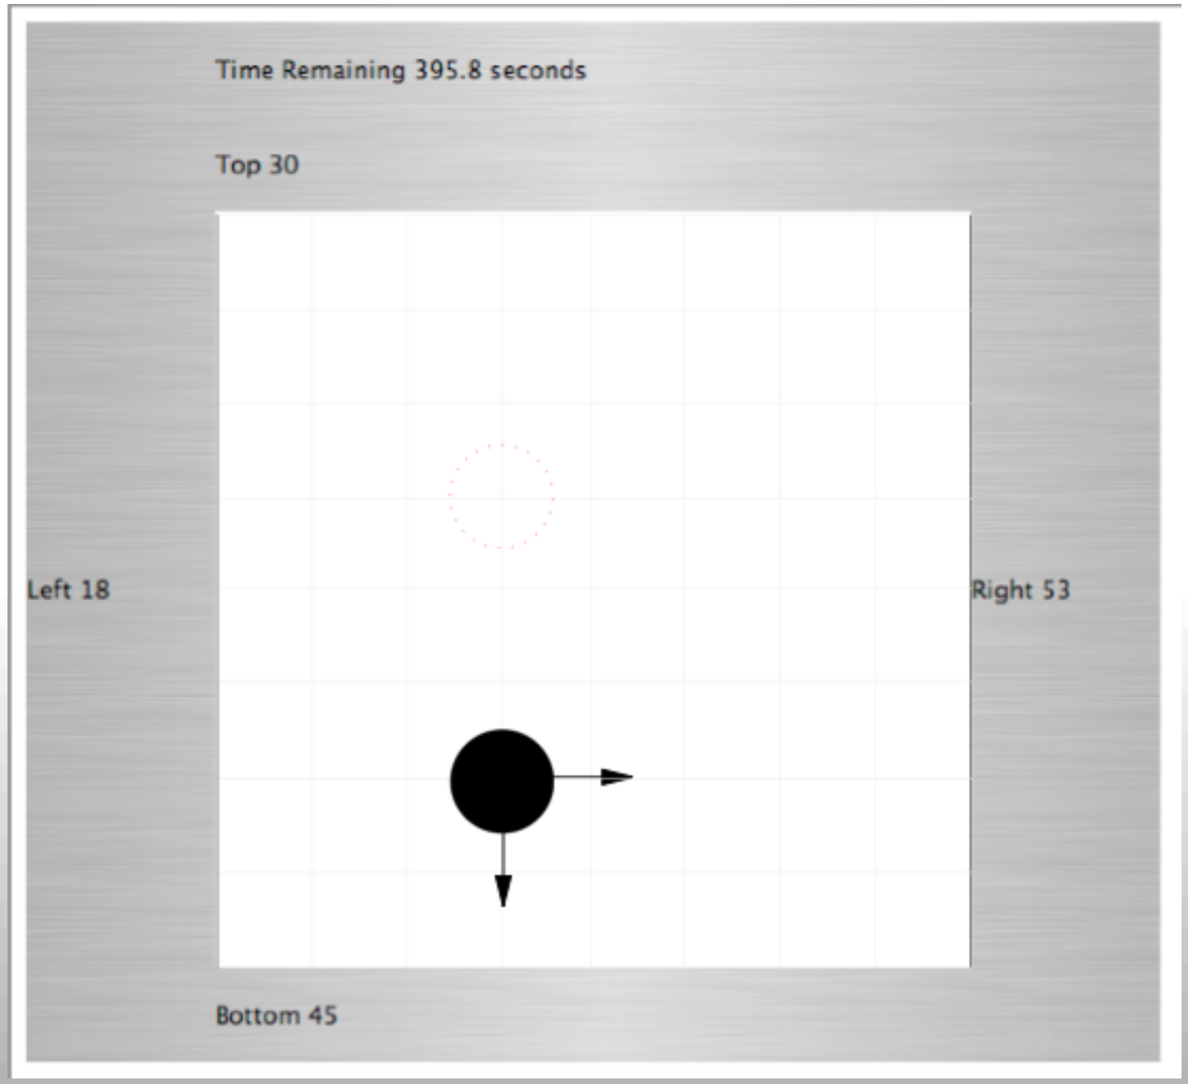

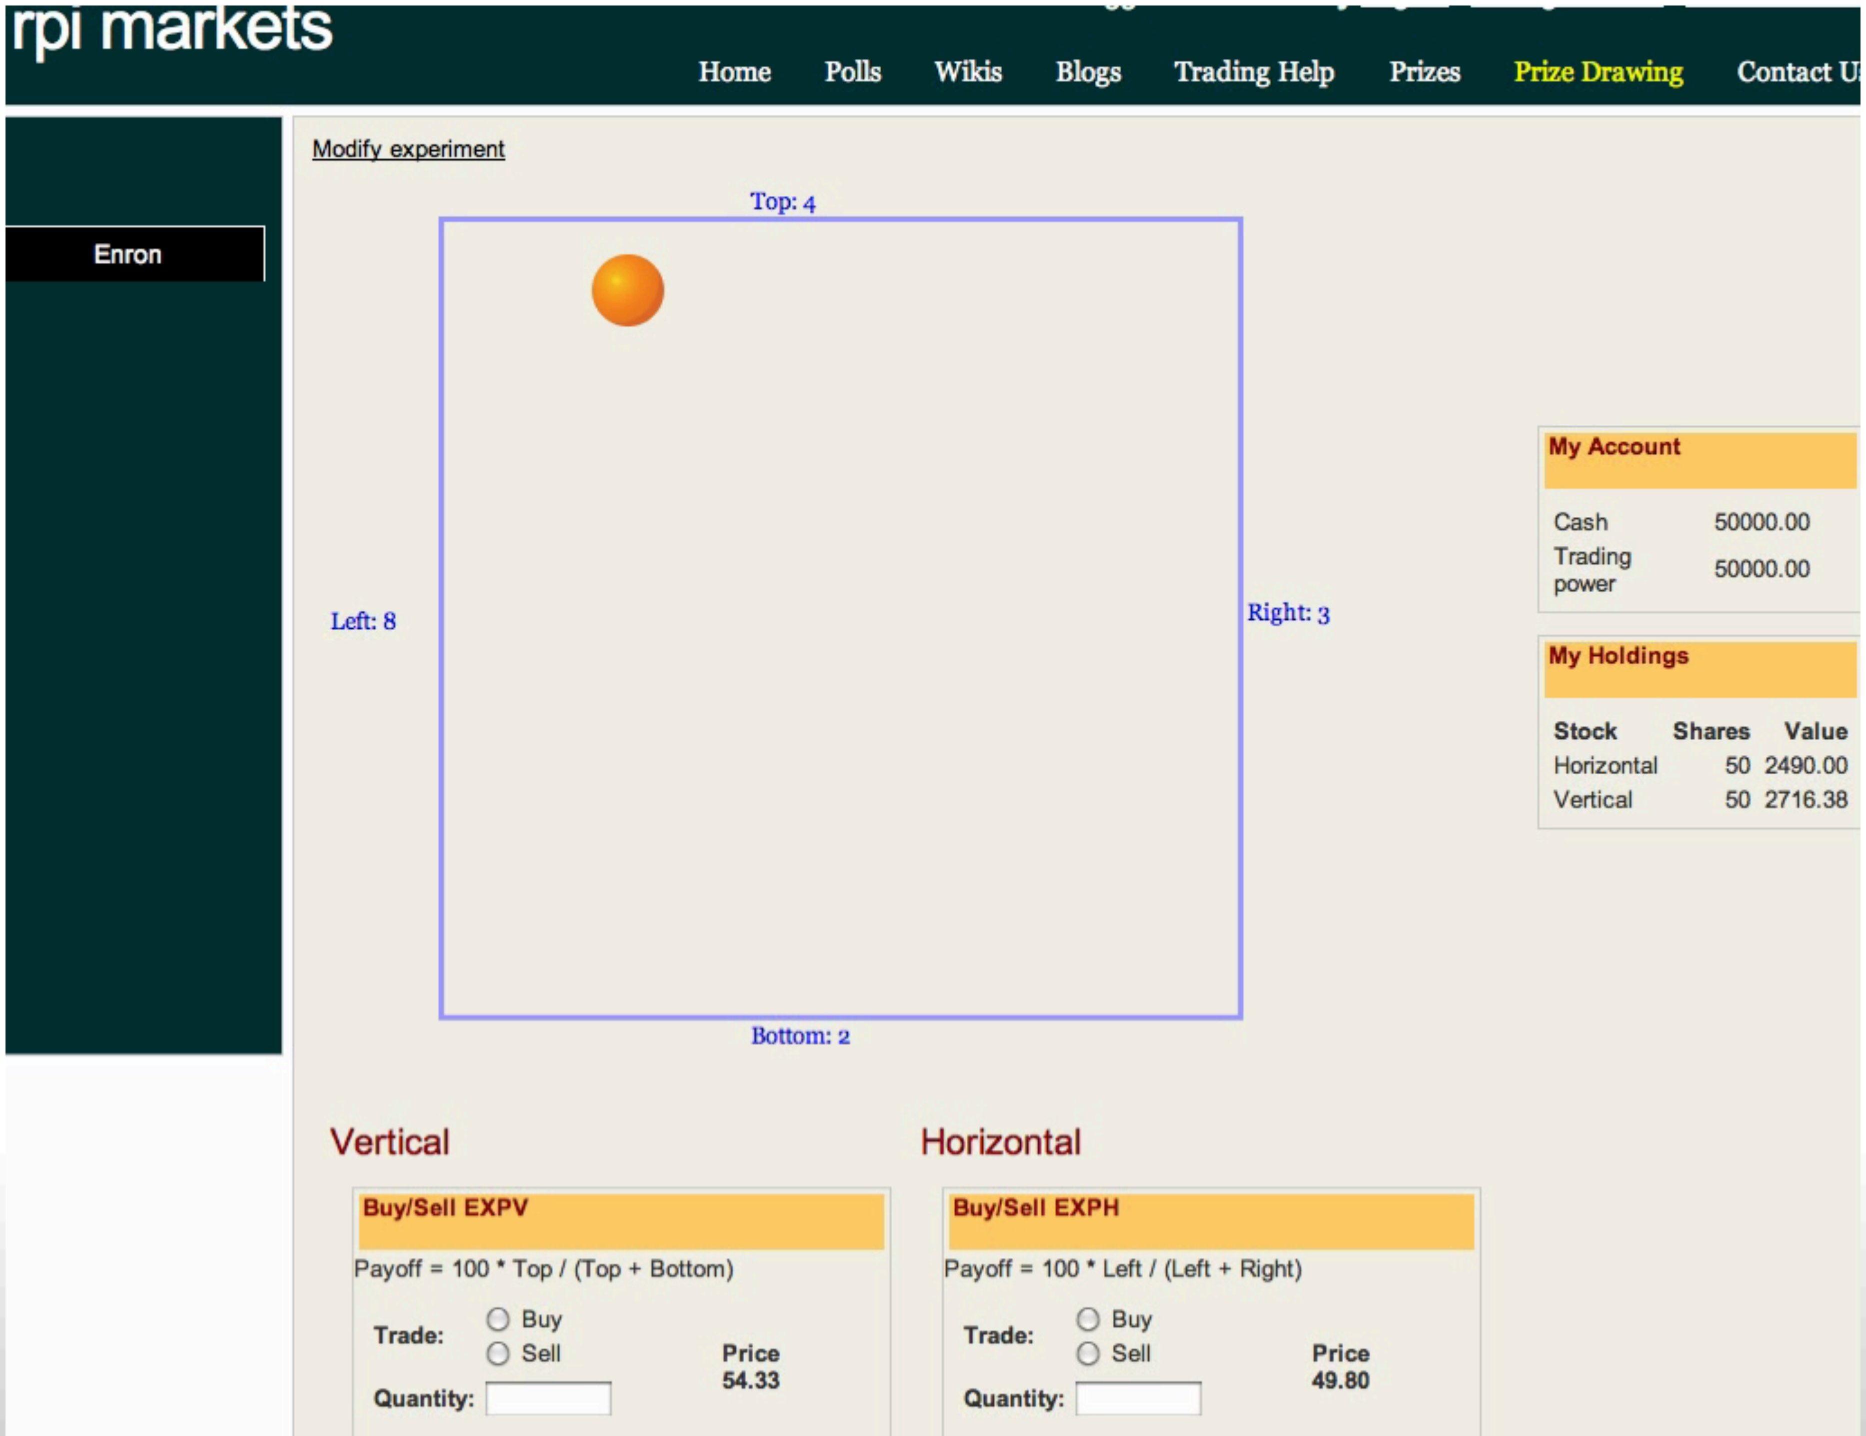

rpi markets
Home
Polls
Wikis
Blogs
Trading Help
Prizes
Prize Drawing
Contact U
Enron
Modify experiment
Top: 4
Left: 8
Right: 3
Bottom: 2
My Account
Cash 50000.00
Trading power 50000.00
My Holdings
Stock Shares Value
Horizontal 50 2490.00
Vertical 50 2716.38
Vertical
Buy/Sell EXPV
Payoff = 100 * Top / (Top + Bottom)
Trade: Buy Sell
Price 54.33
Quantity:
Horizontal
Buy/Sell EXPH
Payoff = 100 * Left / (Left + Right)
Trade: Buy Sell
Price 49.80
Quantity:

Example 1: True Value Constant

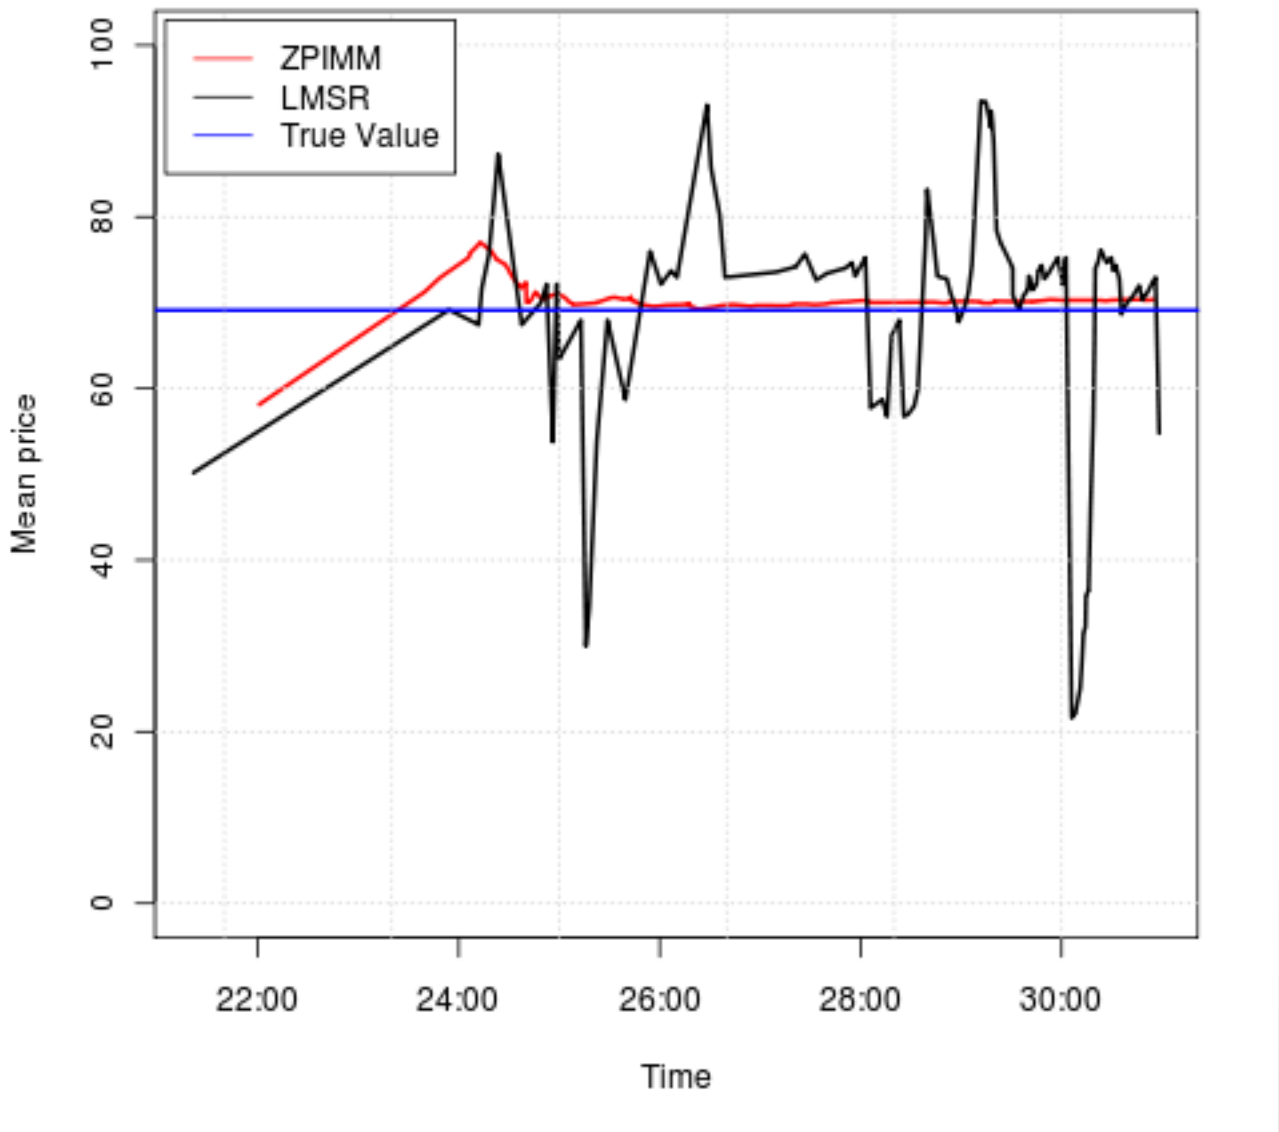

Example 2: True Value Jumps

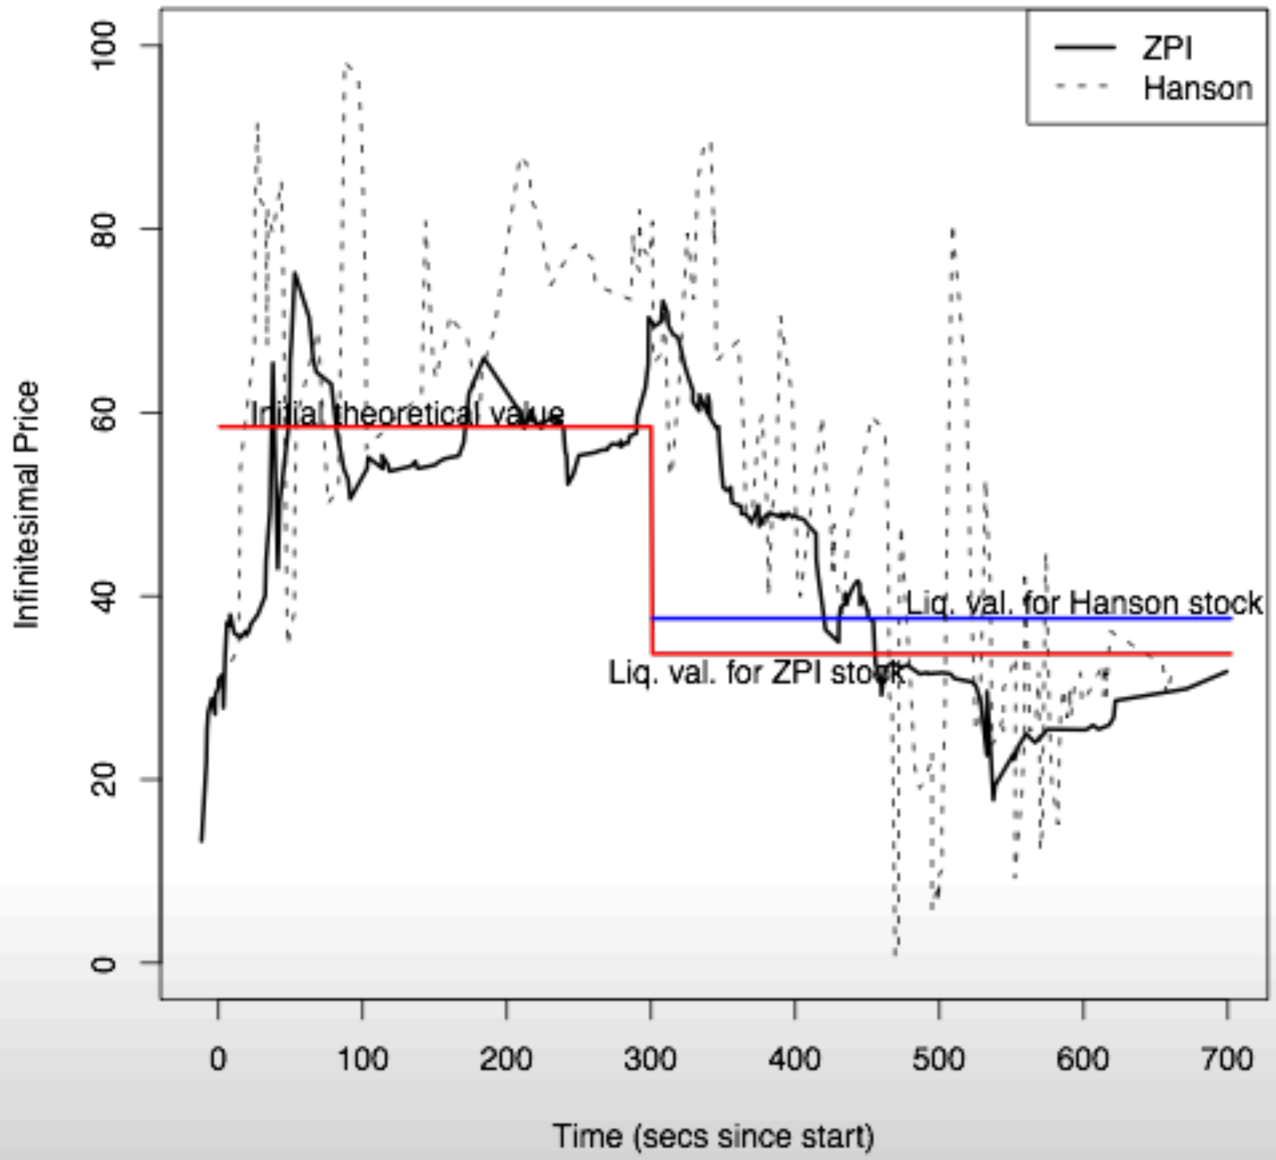

Results

- BMM provides more liquidity and takes less losses (in 5 out of 6 experiments and on average) in doing so than LMSR
- Also provides a more stable price that is closer to the true value

Profits

Expt	BMM	LMSR
1	47231	-1350
2	8973	-1511
3	4084	-1602
4	-10589	-2619
5	9135	-3169
6	35376	1798
6*	20226	-92

Spreads

Expt	BMM	LMSR
1	4.04	3.12
2	3.21	3.06
3	0.74	3.22
4	1.42	3.61
5	0.80	2.84
6	1.33	3.77

Experiment 2: Trading Agents

- Trading bots with access to slowly improving information on coin flip outcomes
  - Simulates “Bouncing Balls”
- Fundamentals Traders
- Learning (“RE”) Traders
- Technical Traders (Moving Average and Range)

	Average profit		Spread		RMSD		RMSDeq	
	bmm	lmsr	bmm	lmsr	bmm	lmsr	bmm	lmsr
10%	-823.74	-1915.51	2.38	2.35	16.09	19.27	5.97	6.63
40%	16630.89	-1496.90	1.24	1.94	12.19	12.95	3.58	6.30
60%	23630.75	-1097.00	1.06	1.88	10.81	14.05	3.10	6.15
100%	-295.61	-3055.04	0.94	1.95	9.28	8.42	3.04	4.87
RE40%	34494.88	-2008.72	1.62	2.02	13.32	14.61	4.87	4.59
RE60%	25223.28	-2312.65	1.28	1.99	11.60	12.05	3.62	4.81
RE100%	-738.83	-3077.43	1.03	1.98	9.67	9.10	3.15	4.56

Experiment 3: Instructor Rating Markets

- Each course has a security liquidating from 0 to 100
  - Market orders through a market making algorithm
- Two-week rating periods
  - Trading accounts start with initial fake money, shares
  - Students in each course rate their instructor
  - Markets liquidate based on this rating
- Prizes
  - 4 Rank-based
  - 1 Rating participation

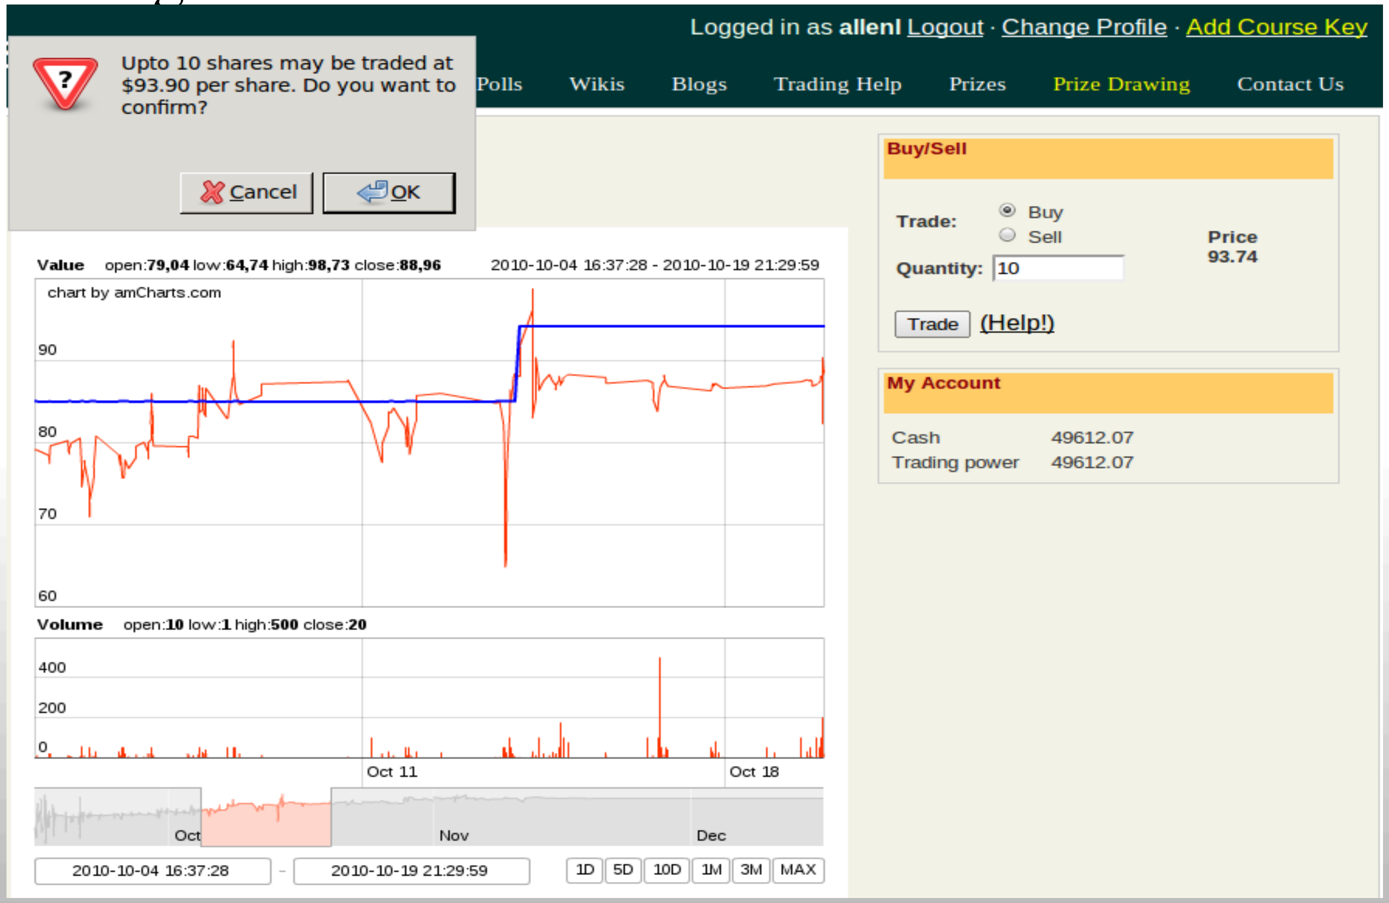

[Chakraborty, Das, Lavoie, Magdon-Ismail & Naamad, AAAI 2013]

Experiment 3 Results

MM	#Periods	Avg profit	Max loss	Std. dev. of prices	Dev. from next liquidation
LMSR	25	1338.99	-5298.58	8.6	16.9
BMM	15	8273.13	-13763.40	3.0	9.6

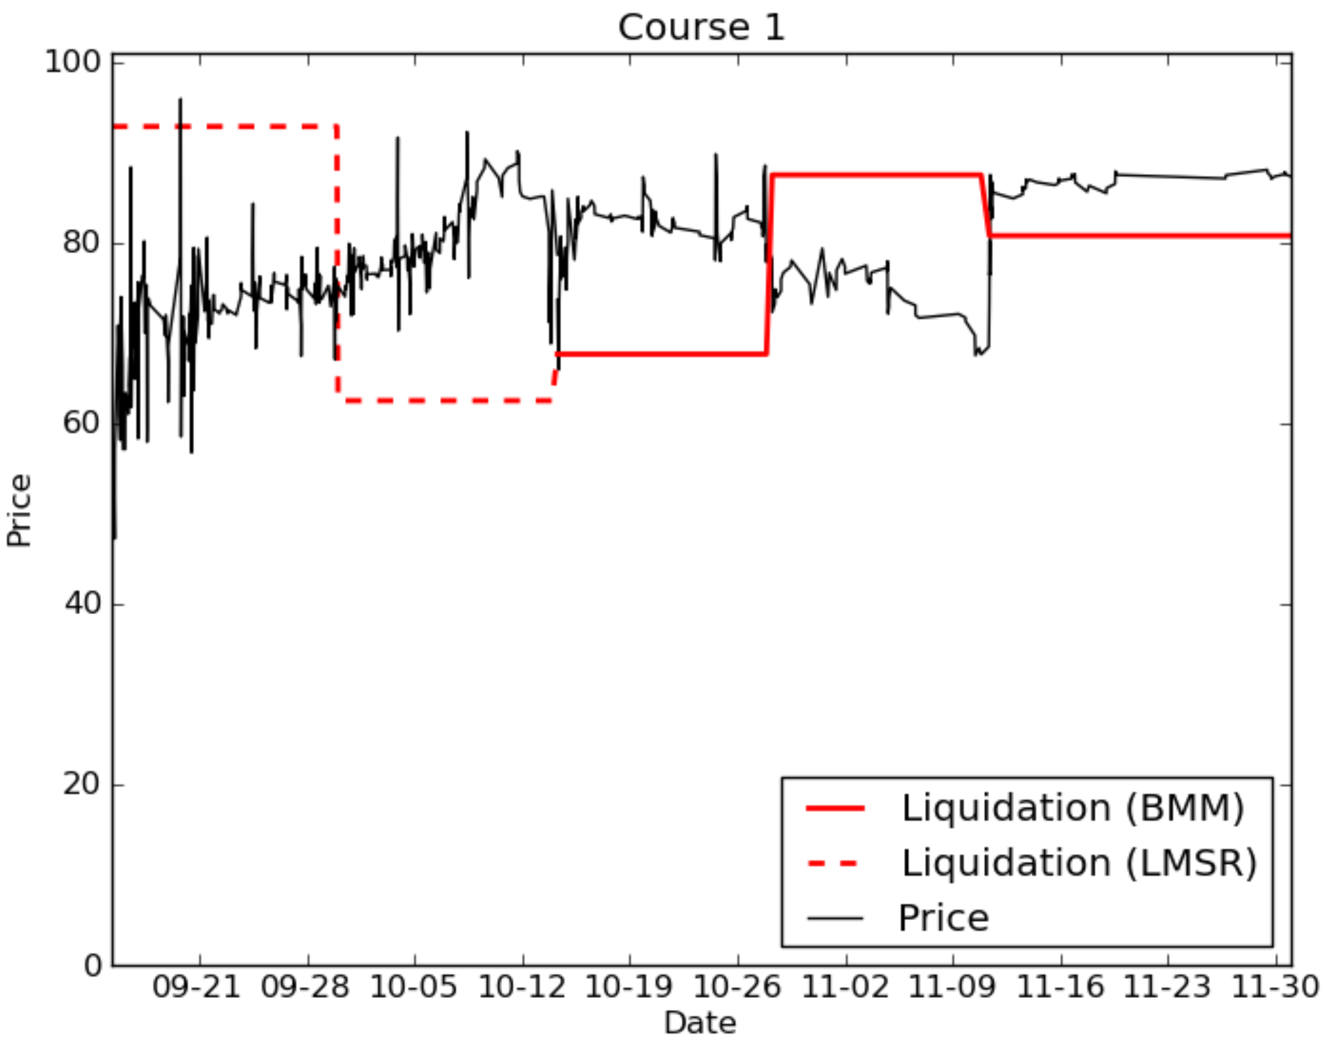

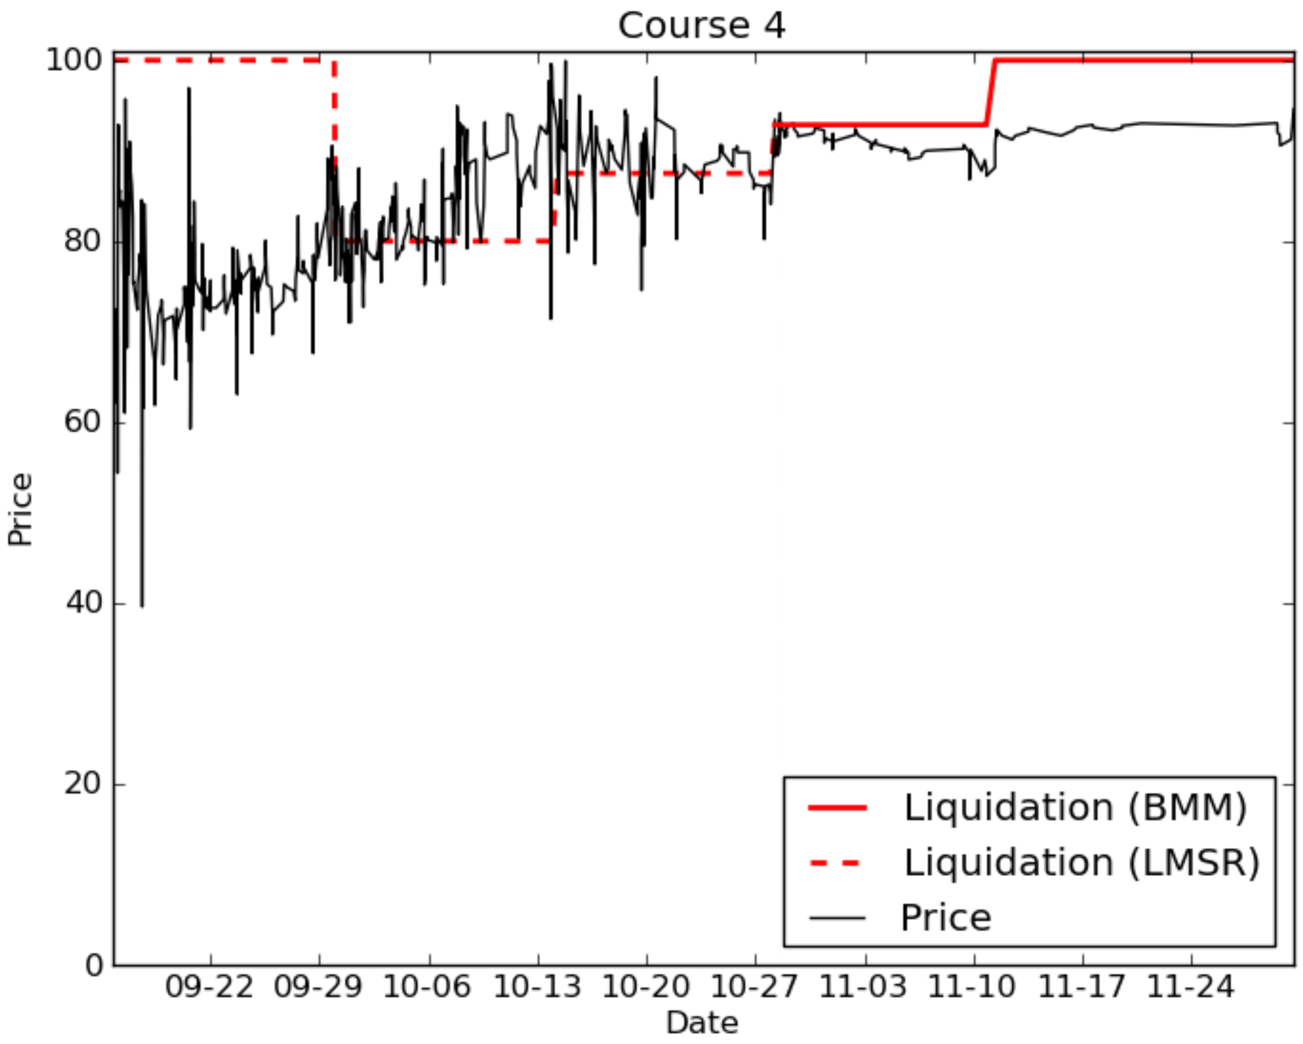

Market outcomes

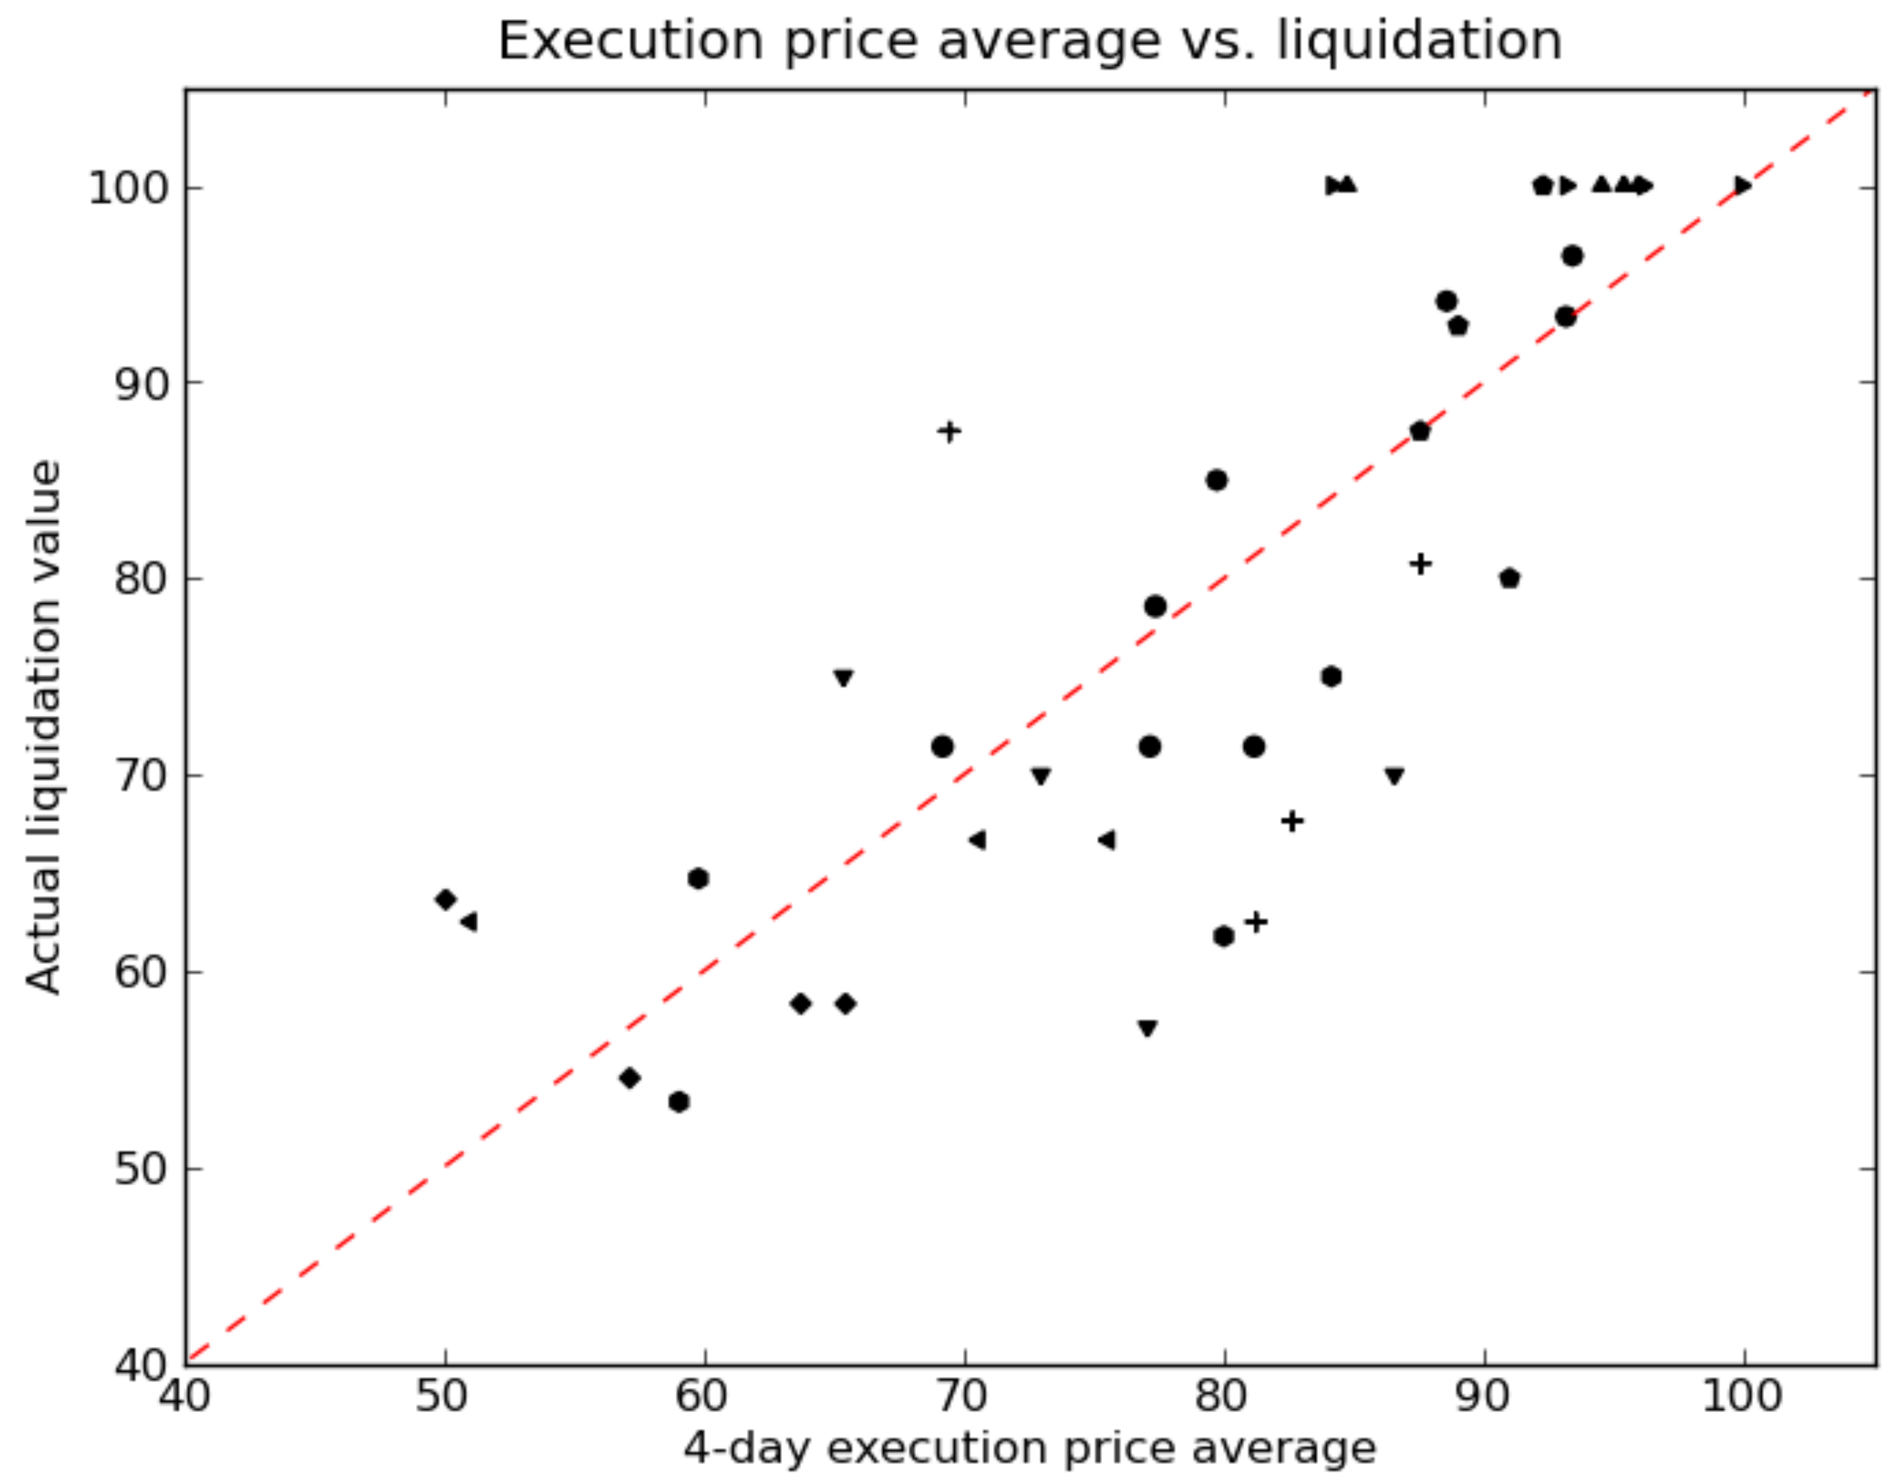

Prices incorporate new information

- Linear model predicting future liquidations
  - Previous liquidation
  - Market price average

- Price average is more predictive
  - R-squared (0.58 vs. 0.48)
  - Previous liquidations insignificant in linear model

$$\text{Liq}_{s,\rho} = \beta_1 \text{Liq}_{s,\rho-1} + \beta_2 \text{Price}_{s,\rho} + \alpha$$

α est.	β_1 est.	β_2 est.	Sample Size
7.02	0.17	0.72 **	40
		$^{**}p < 0.01$	

Raters provide new information

We know which traders are raters for a class ("in class").

How do we tell the informational difference between "in class" and "out of class" traders?

Examine trades that originate at prices in-between previous and future liquidations

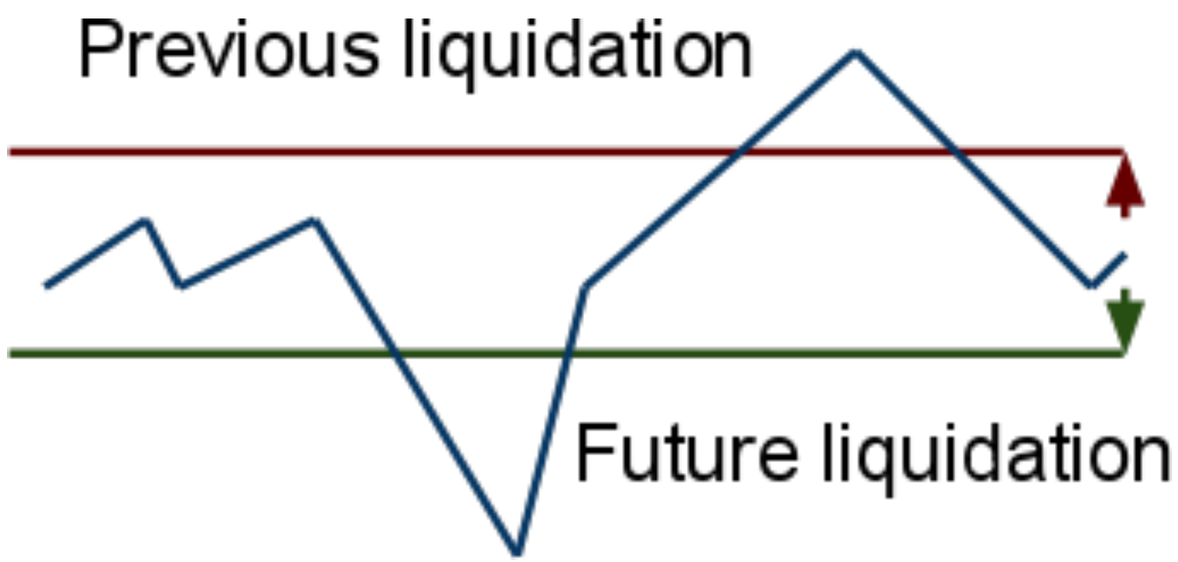

- Results:
  - In-class traders: toward future liquidations 54% of the time (significantly > 50%)
  - Out of class traders: toward future liquidations only 48% of the time (significantly < 50%)

IRM ratings are real

- Correlation of IRM and official department evaluations
  - Seven computer science courses
  - Last three trading periods
  - Ratings 0.86
  - Prices 0.75

- Prices predict ratings, ratings predict evaluations despite:
  - Smaller sample sizes
  - Manipulation potential