Shopbot Economics

Jeff Kephart and Amy Greenwald

IBM Institute for Advanced Commerce
IBM Thomas J. Watson Research Center
Hawthorne, New York

July 6, 1999

Economics of Information

George Stigler [1961]

- price dispersion is attributed to costly search procedures

Shopbots Today

- shopbots specialize in collecting and distributing price information at low cost

Pricebots Tomorrow

- automated agents that set prices in attempt to maximize profits for sellers, just as shopbots seek to minimize costs for buyers

Adaptive Pricebot Strategy

Cyclical price wars arise when pricebots are myoptimal (*i.e.*, myopically optimal) and $0 < w_1 < 1$.

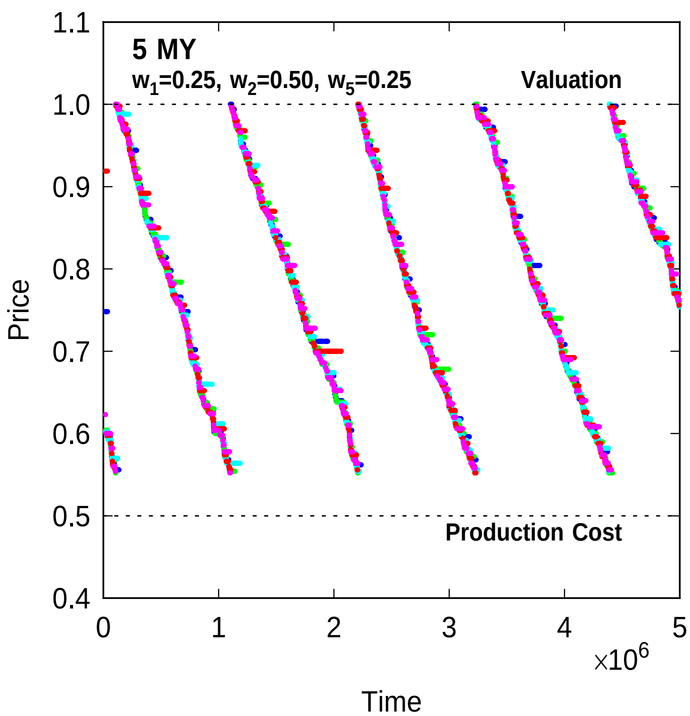

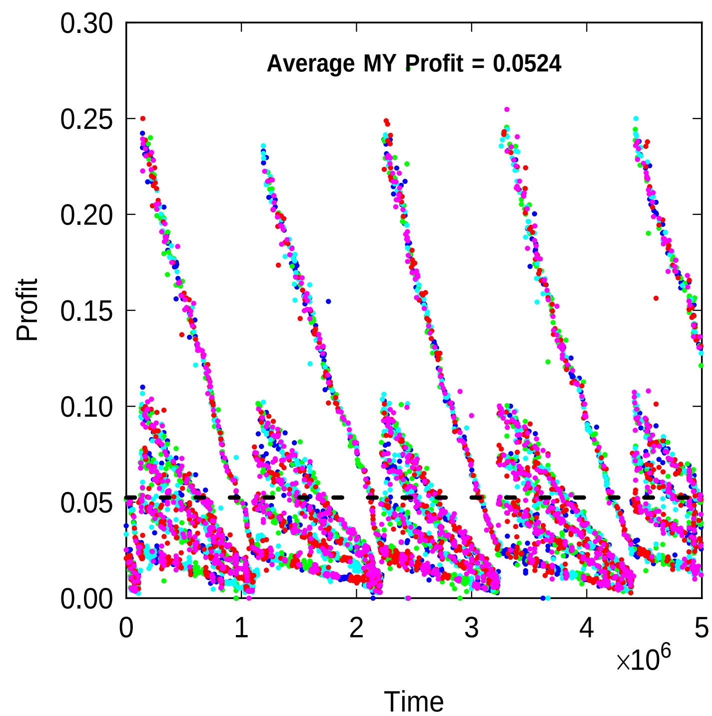

There is no pure strategy Nash equilibrium for pricebots.

Game-Theoretic Pricebot Strategy

Mixed strategy Nash equilibrium Pricebots choose prices at random according to probability distribution $f(p)$.

5 pricebots, $w_1 = 0.2$, $w_2 + w_5 = 0.8$.

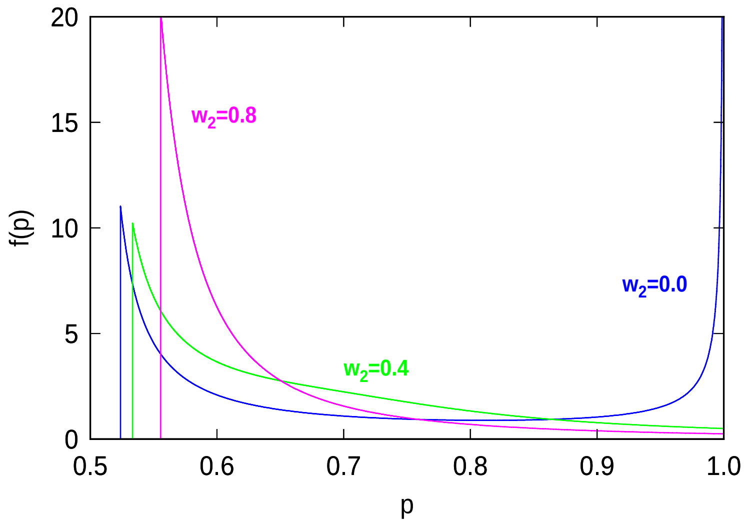

20 pricebots, $w_1 = 0.2$, $w_2 + w_{20} = 0.8$.

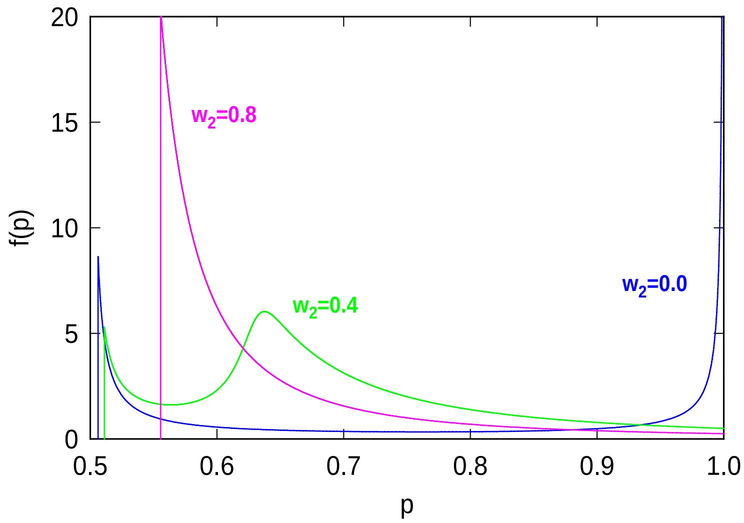

Adaptive pricebots learn Nash equilibrium via regret algorithms *e.g.*, Foster and Vohra, 1997.

Rational Buyer Strategy

Total Buyer Expenditure = Expected Price + Search Costs

Buyer Price Distributions

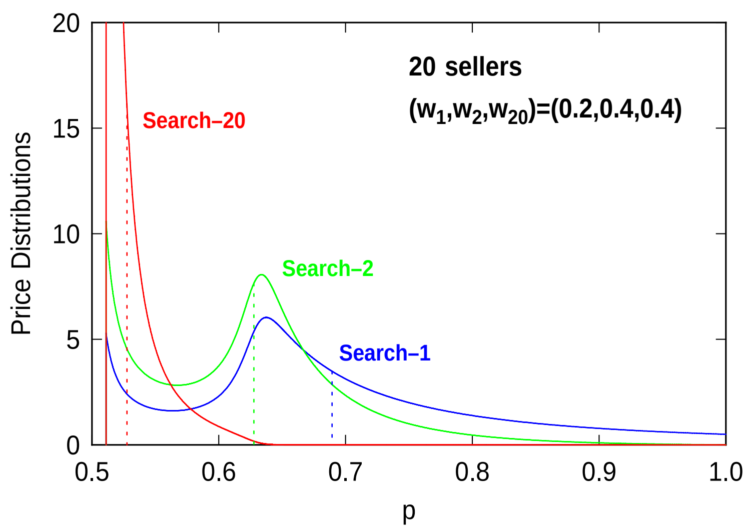

Average Buyer Prices

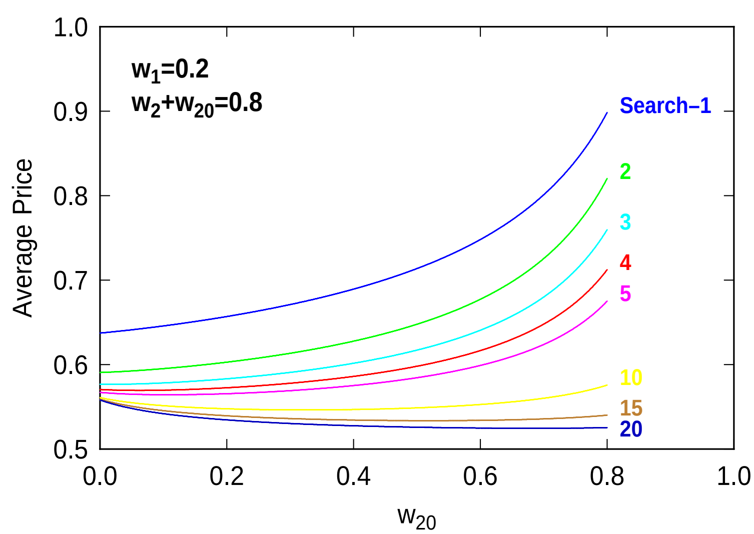

Value of Information = Expected Savings - Search Costs

Adaptive Buyer Strategy

At each time t

1. Small fraction of buyers switch from their present search strategy to current optimum.

2. Sellers compute new game-theoretic pricing strategy.

Linear Search Costs

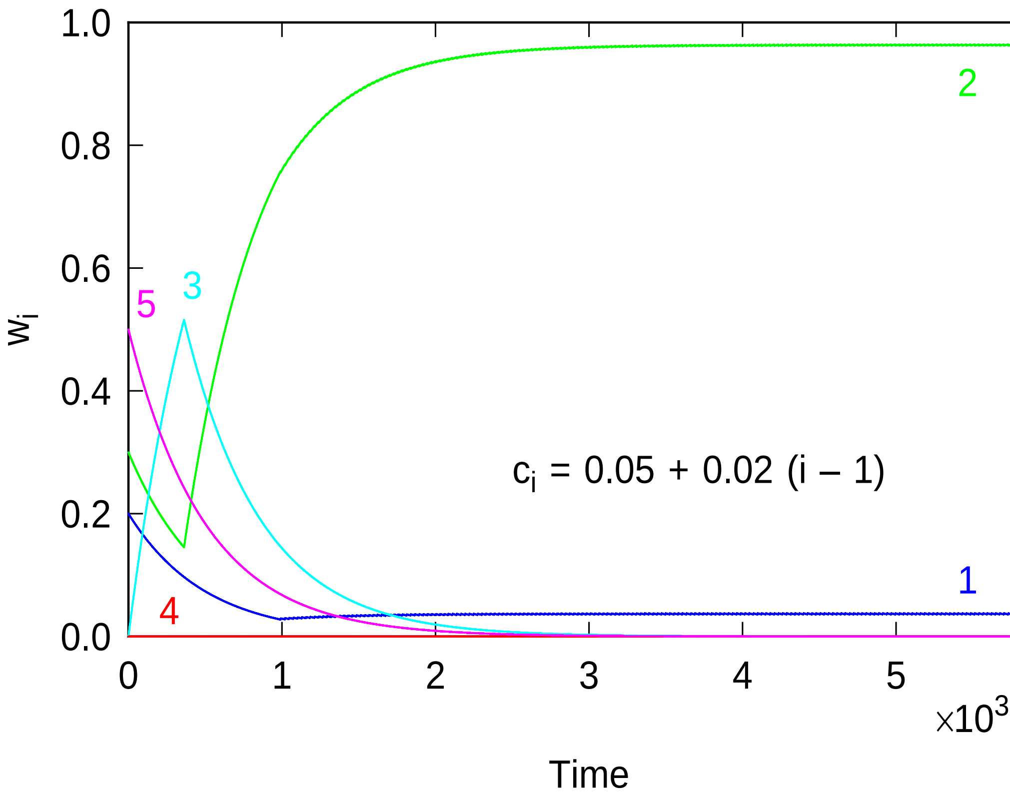

$$\begin{aligned} \text{Initial state:} \quad (w_1, w_2, w_5) &= (0.2000, 0.4000, 0.4000). \\ \text{Final state:} \quad (w_1, w_2, w_5) &= (0.0141, 0.9859, 0.0000). \end{aligned}$$

Burdett and Judd, 1983 Linear search costs yield $w_1 + w_2 = 1$.

Adaptive Buyer Strategy

Shopbots drastically lower search costs Assume costs are non-linear in the number of searches.

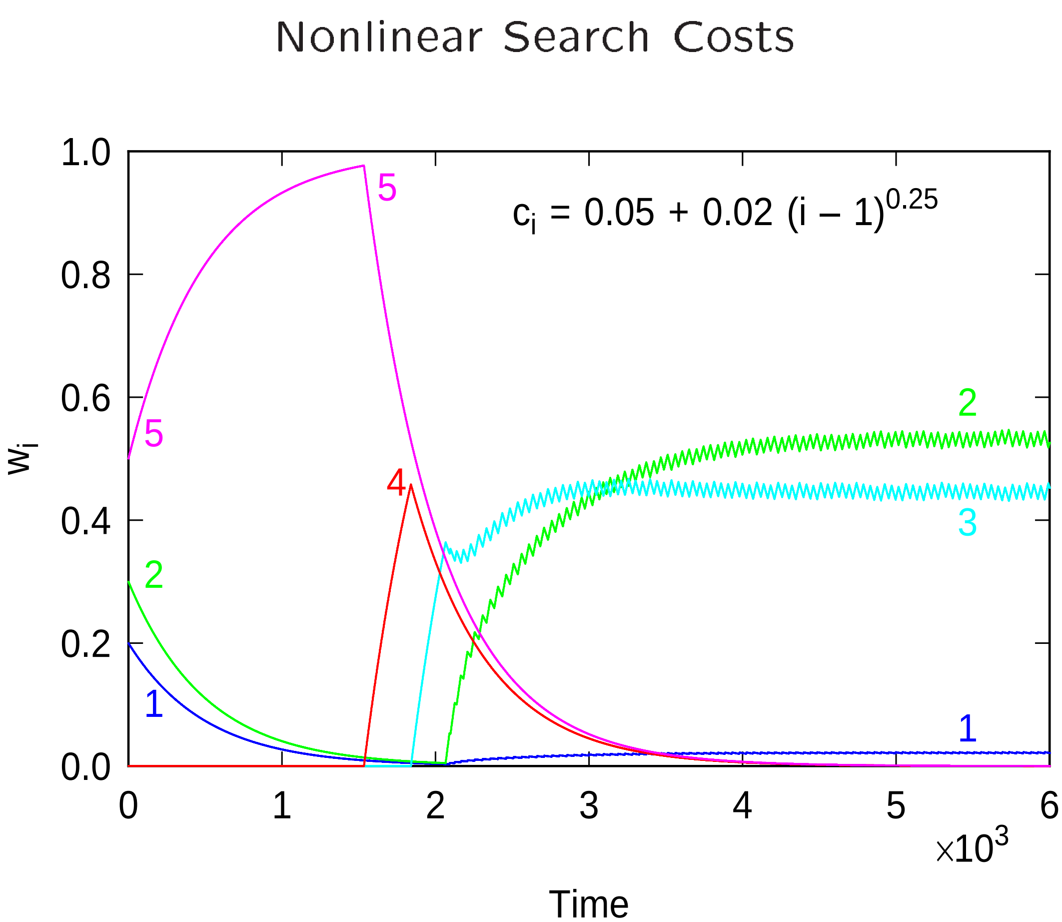

$$\begin{aligned} \text{Initial state:} \quad &= \quad (0.200, 0.300, 0.000, 0.000, 0.500). \\ \text{Final state:} \quad &= \quad (0.020, 0.550, 0.430, 0.000, 0.000). \end{aligned}$$

Nonlinear search costs can yield more complex, even chaotic, mixtures of strategies.

Fixed + Adaptive Buyers

Suppose small fraction of buyers fixate on search-1, regardless of what strategy is optimal, while other buyers adapt.

4% Fixed Search-1 Buyers

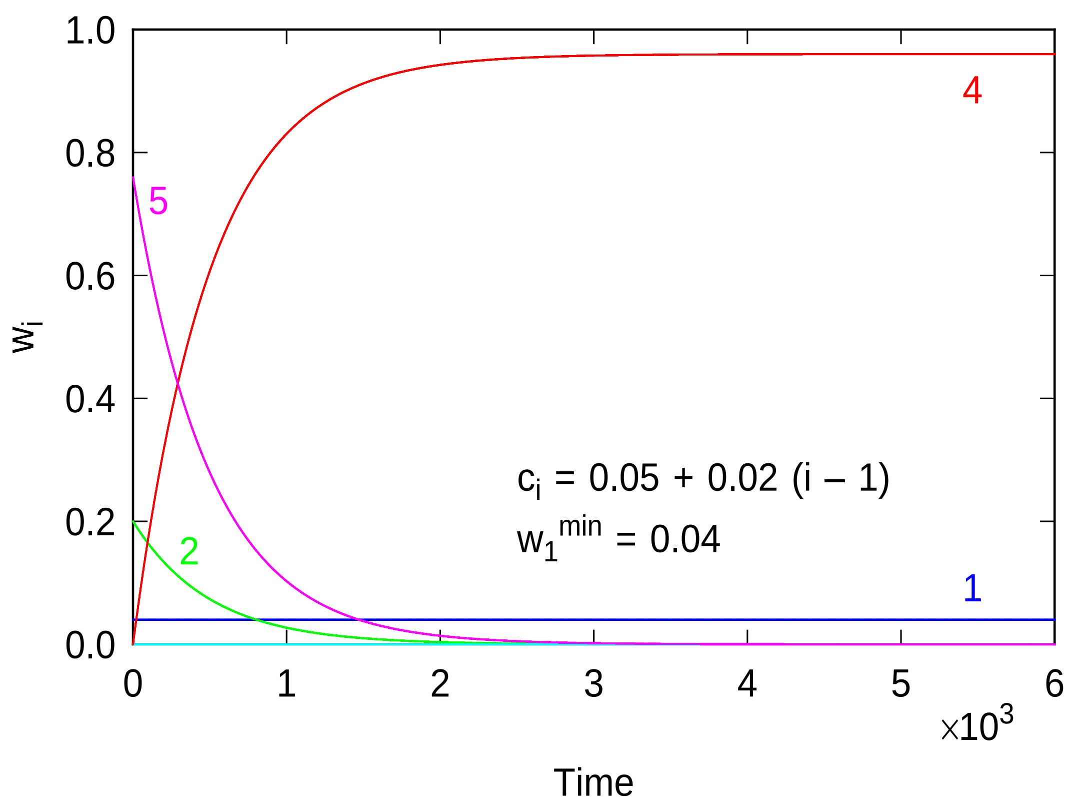

$$\begin{aligned}\text{Initial state:} \quad &= \quad (0.040, 0.200, 0.000, 0.000, 0.760).\\ \text{Final state:} \quad &= \quad (0.040, 0.000, 0.000, 0.960, 0.000).\end{aligned}$$

Mixture of fixed and adaptive buyer behavior can lead to strategies other than just search-2 co-existing with search-1.

Fixed + Adaptive Buyers

When some buyers fixate on search-1 strategy, the others adapt co-existing strategy depending on initial conditions.

Basins of Attraction

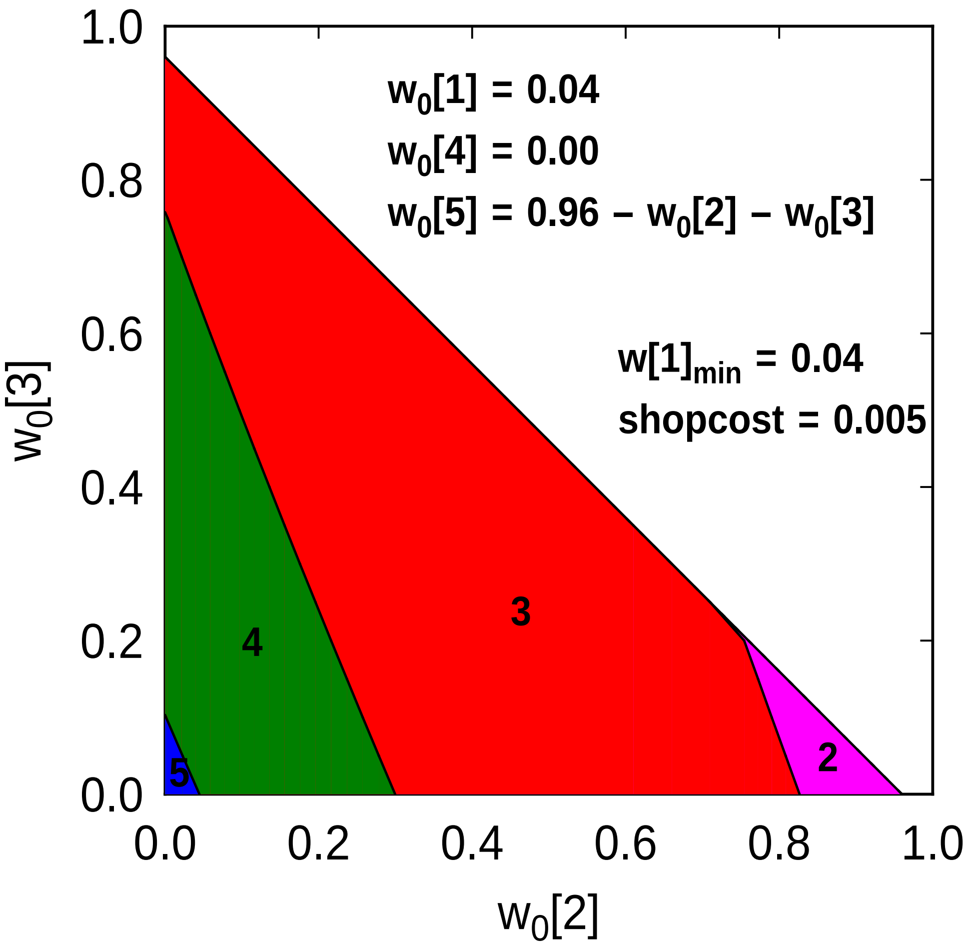

$\lfloor w_1 \rfloor$	$\delta = 0.001$	$\delta = 0.005$	$\delta = 0.020$
0.01	5	2	2
0.04	5	2–5	2
0.20	5	5	2–3

Fixed vs. Adaptive Pricebots

Introduce 1 MY pricebot into group of 4 GT pricebots . . .

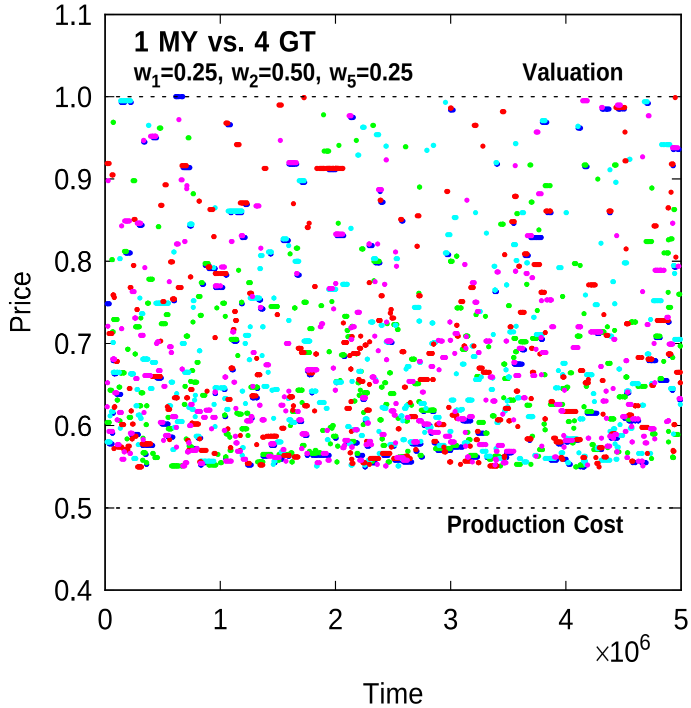

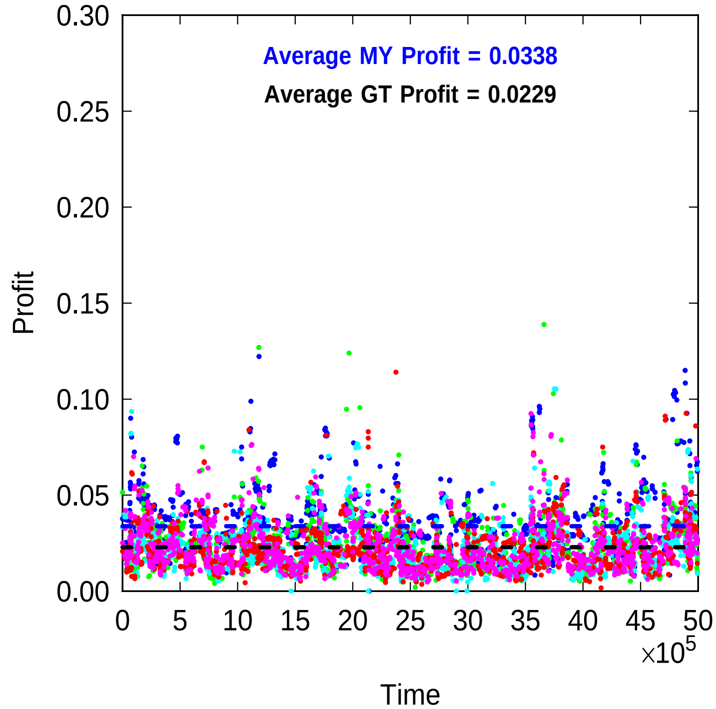

. . . MY outperforms GT.

Informed, Adaptive Pricebots

5 MY Pricebots . . .

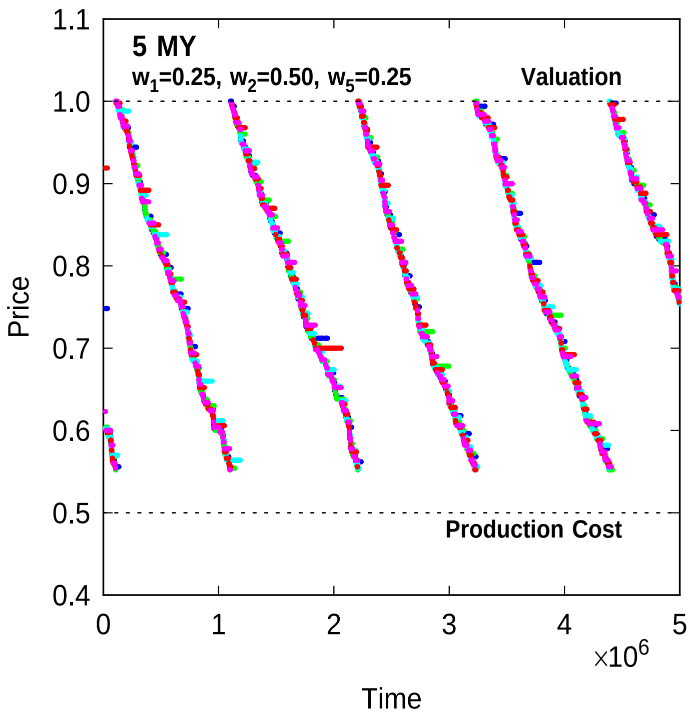

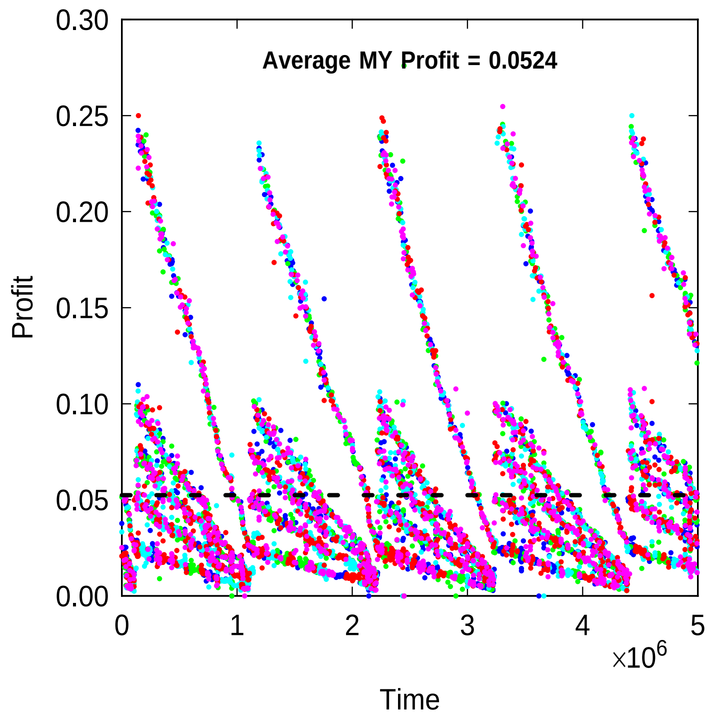

- MY profits (0.0524) more than twice GT profits (0.025);
- but instabilities in the form of cyclical price wars arise;
- and MY pricebot requires knowledge of buyer demand and other sellers' prices, which may be costly to obtain.

Naive, Adaptive Pricebots

Derivative-following (DF) Strategy Adjust price in same direction as long as profit increases; otherwise reverse the direction of price adjustment.

5 DF Pricebots . . .

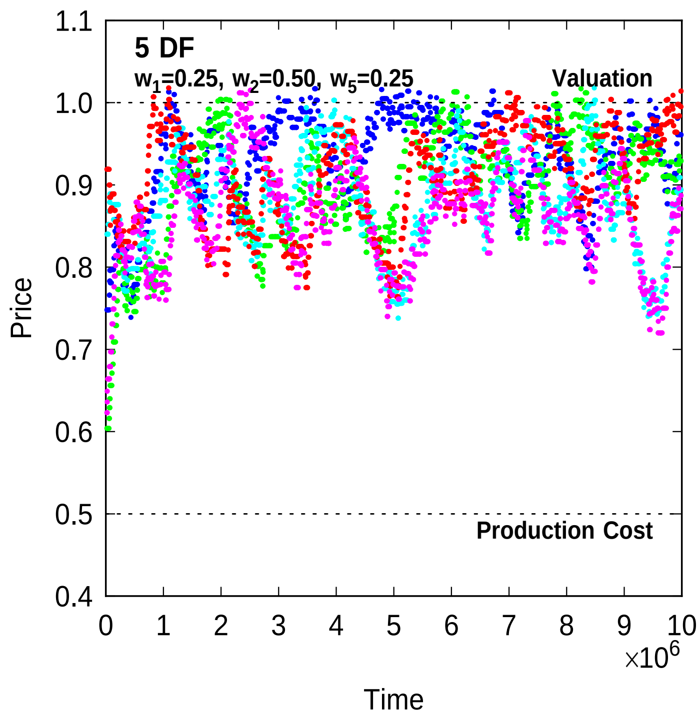

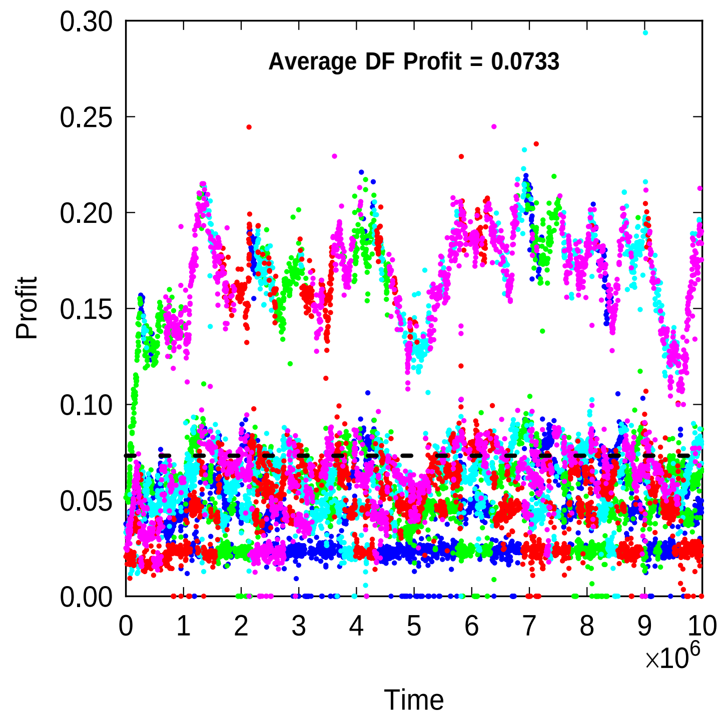

- Tacit collusion results: *i.e.*, an effective cartel despite no actual communication!

- Average profit is nearly 3 times that of GT pricebots. Perfect cartel would achieve profit of 0.1 per pricebot.

- Requires no knowledge of sellers' prices or buyer demand; price-setting mechanism based on historical observations.

Informed vs. Naive Pricebots

Introduce 1 MY pricebot into group of 4 DF pricebots . . .

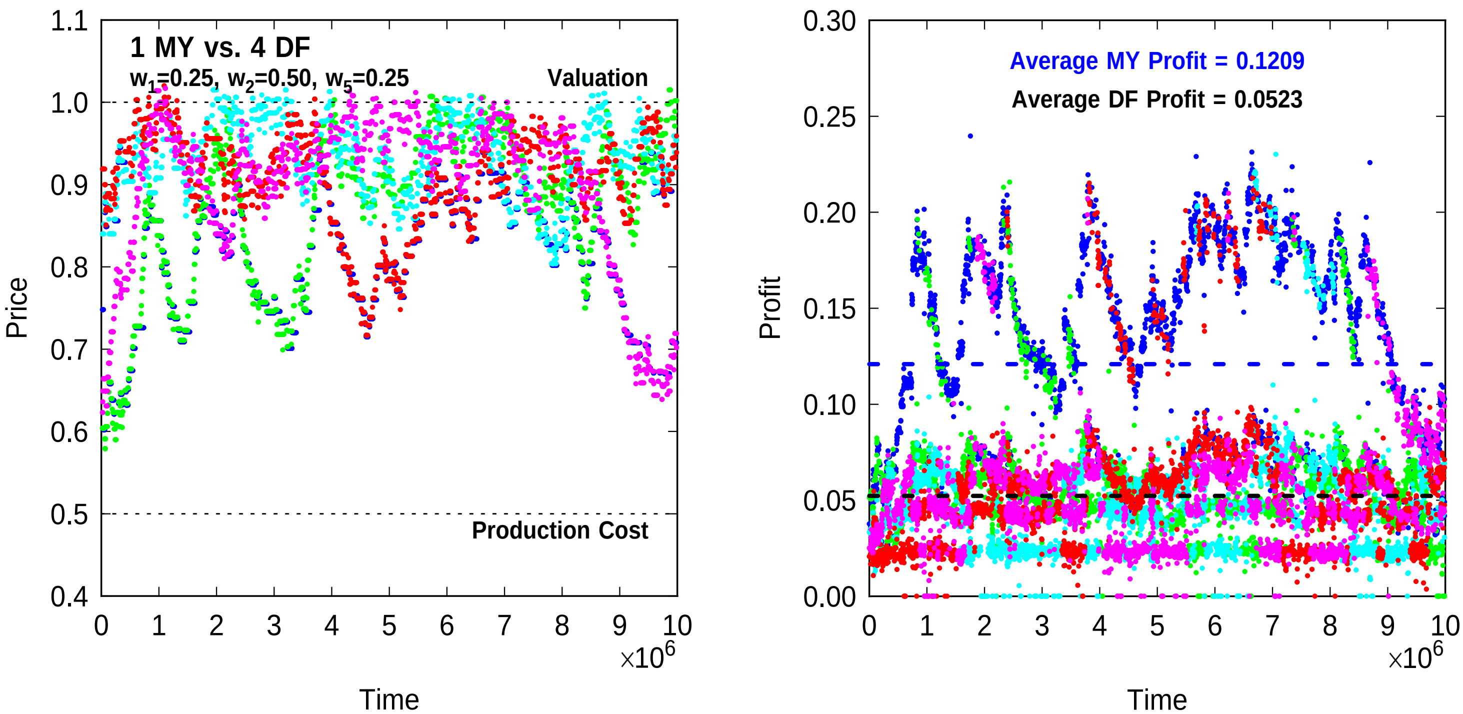

. . . and it will exploit them mercilessly, stealing their profits, and earning more than twice what they do!

Q-Learning Pricebots

Watkins, 1989 Reinforcement Learning Scheme

2 MY Pricebots . . .

2 Q Pricebots . . .

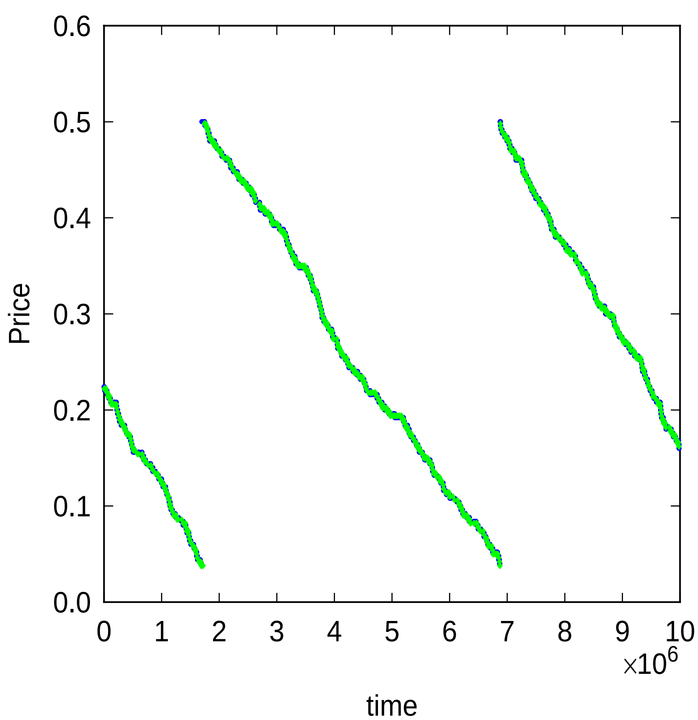

1MY vs. 1MY

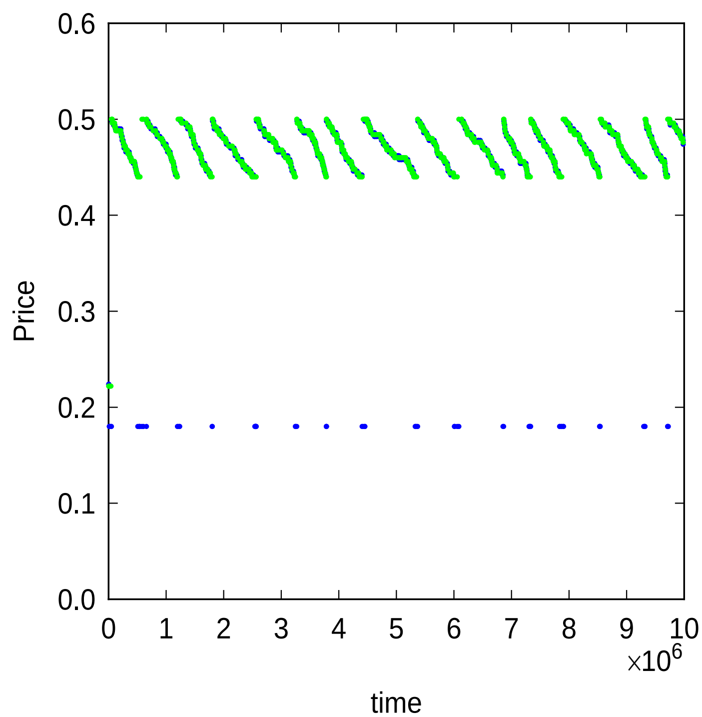

1Q(a) vs. 1Q(b)

- Q pricebots detect and abandon price wars early on
- Q profits (0.125, 0.117) exceed MY profits (0.089, 0.089)

No External Regret Pricebots

Freund and Schapire, 1995 Probabilistic Updating Scheme

2 NER Pricebots . . .

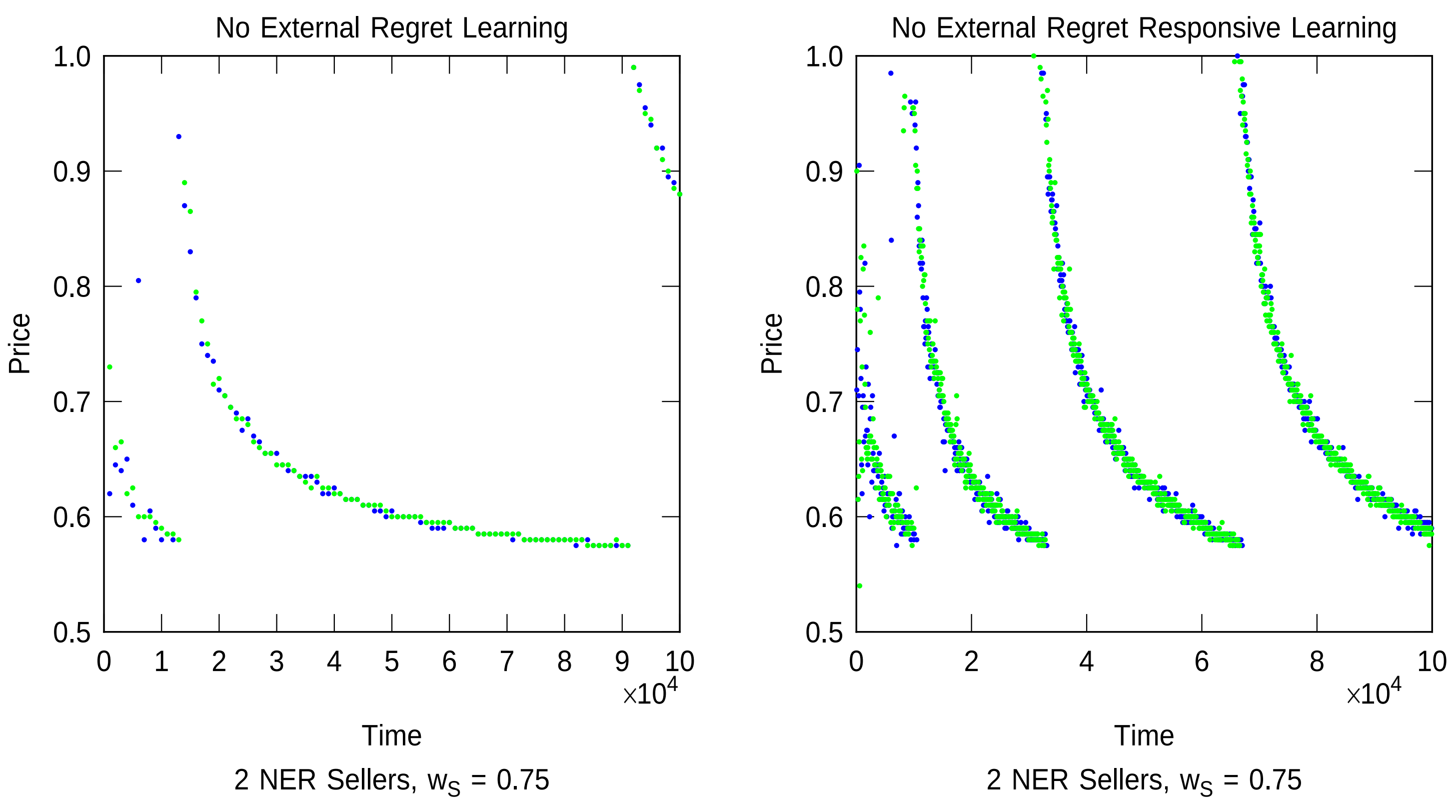

- NER pricebots cycle through prices exponentially
- responsive NER pricebots engage in limited price wars

No Internal Regret Pricebots

Foster and Vohra, 1997 Converge to Correlated Equilibrium

2 NIR Pricebots . . .

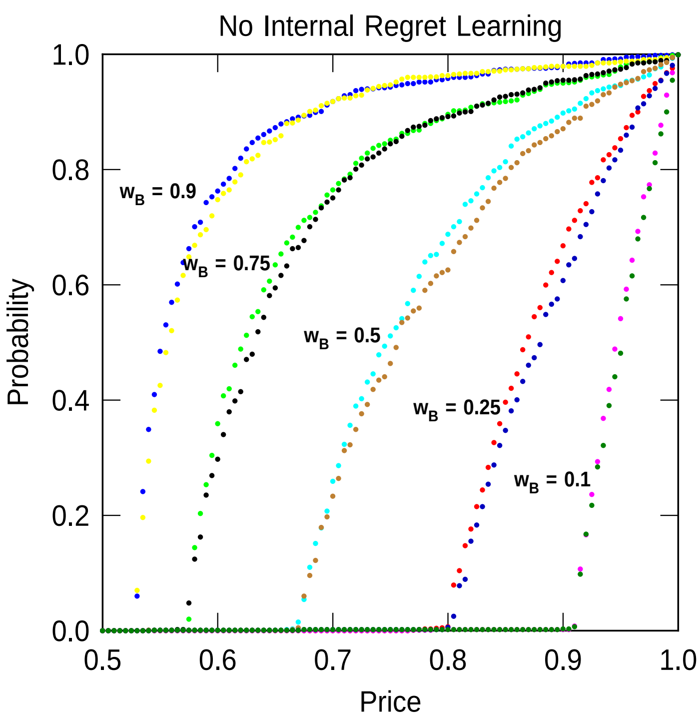

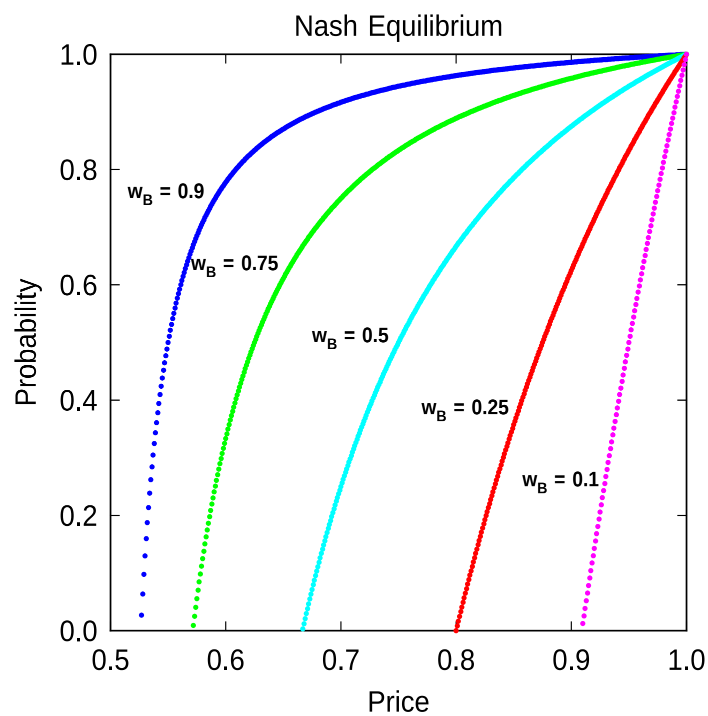

. . . NIR pricebots learn Nash equilibrium!

Future Work

- Study dynamic interplay of adaptive buyers and sellers
- Effects of multiple attribute product differentiation
- Dynamic pricing of price and product information in full-fledged economy of software agents, consisting of buyers, sellers, and economically motivated shopbots

Shopbot Economics forms part of the Information Economies project at IBM Research's Institute for Advanced Commerce. The project goal is to:

> accurately describe and predict collective interactions of billions of economically motivated software agents, and use insights so gained to design agent strategies, protocols, and infrastructures.

Project description and research papers available at:

www.research.ibm.com/infoecon