Market Simulation

Social Intelligence

Daniel Polani

Market Models

Motivation:

- markets require exchange of commodities and information
- internet provides a fast and reliable information-exchange mechanism
- information exchange may be translated into commodity exchange via trustworthy partners

Trustworthyness

Question: Where does it come from?

Reminder: consider Prisoner's Dilemma

- Nash-equilibrium is Pareto-dominated; i.e. the stable ("paranoid") equilibrium is not the overall most desirable
- can be avoided by iterative play
- i.e. knowing that one has to interact with another agent again is an incentive for "trust" and to move away from "paranoid" Nash-equilibrium
- agents can form a "personalized" model of their partner's trustworthyness

Trustworthyness II

Note: in large societies with a high degree of anonymity of interactions, trustworthyness is established via norms (see [Flentge et al. 2003] for an example and references)

In Reality: trustworthyness is established by interface agents

- government (currencies)
- large companies (e.g. credit cards)

Remark: trustworthyness allows converting information int commodities (banks)

Trustworthyness and Internet

Observation: trustworthyness is often established through physical properties (money); in earlier times information exchange and physical contact was more intimately related

Internet: establishes complete physical detachment while allowing information exchange

Bottom Line: internet requires specific mechanisms to establish trustworthyness.

For now: concentrate on information exchange and ignore issues of trustworthyness; we assume them as given

Market Modeling

Consider Model:

- market as server
- agents (buyers/sellers) as clients
- agents are either buyers or sellers, not both
- agents can choose to send out and to ask for offer
- only one type of commodity
- retail market (no grossists)

Price Establishment Models: choice between

- bazaar (negotiate on price for each transaction)
- supermarket (buy it or leave it)

DMarks Model

Model: as before, choosing

- supermarket model of price establishment
- luxury goods (no minimal consumption necessary)
- asynchronous update

Sellers and Buyers

Sellers:

- small start capital
- flexible setting of production
- storage costs
- production costs
- **goal:** profitability

Buyers:

- fixed income per time
- two types: **A**nybuyers and **B**argainhunters
- **goal:** consumption, while maintaining economical rationale

Transaction: Details

Goods: single, non-perishable commodity traded against currency

Time: discrete time steps

Sellers: in each time step, sellers produce and store real-valued quantities of the good and fix a price for which it will be sold during this step.

Buyers:

- query the current price from a seller and make a one-shot decision from which seller to buy.
- try to purchase good according to demand function from the chosen seller
- if served, consume and receive a reward

Simulation Cycle

- for each time step
  - buyers reset their demands
  - sellers pay store costs, set production levels and produce
  - sellers set new price
  - buyers obtain price quotes and select seller to buy from
  - buyers try to purchase goods from the seller chose
  - buyers consume goods and receive rewards
  - if seller account negative, remove it

Seller Behaviour

Model:

1. sellers have account initialized with small credit c at the start of simulation
2. keep track of the inventory S of their good
3. S incurs store costs c_s per unit per timestep
4. when producing, sellers are charged production costs c_p per unit
5. profit will be placed on the account
6. if account depleted, seller removed from simulation

Seller Simulation

1. initialize account: $A \leftarrow c$
2. pay store costs: $A \leftarrow A - Sc_s$
3. set production level to P
4. produce: $S \leftarrow S + P$, $A \leftarrow A - Pc_p$
5. set price to p
6. serve customer purchases q_i until next timestep: $A \leftarrow A + pq_i$
7. if not broke, goto 2, else quit

Key Steps: setting P and p

Buyer Behaviour

Model:

1. buyers keep account A,
2. on which they receive a fixed income i each timestep.
3. A determines the budget available to satisfy their demands
4. buyers get informed about current market prices by querying sellers
5. obtaining price quotes may be costly for buyers; therefore distinguish two types of buyers:
   (a) *Anybuyers* randomly select one of the sellers to query every timestep
   (b) *Bargain hunters* query all the sellers and can compare prices
6. consider parameter: proportion $\omega \in [0, 1]$ of fraction of fully informed buyers
7. with increasing ω competition between sellers increases, induced by higher price transparency

Buyer Demand

Demand: $d(p)$

Model:

$$d(p) = \max\left\{0, d_{\text{max}} - p\frac{d_{\text{max}}}{p_{\text{max}}}\right\}, \qquad p \geq 0$$

Timestep: demand is reduced after consuming goods during timestep, and reset at beginning of next timeste

Buyer Reward

Reward: need a model to reward consumption, but, at the same time to encourage economical behaviour

Ansatz: $r(p, q)$ determined as the difference between what they would have been willing to pay for the goods according to their bid level and what they actually paid:

$$r(p, q) = \int_0^q d^{-1}(s)ds - pq$$

Buyer Simulation

1. receive income: $A \leftarrow A + i$
2. query one or all sellers (according to buyer type) for prices p_i
3. select seller to purchase from: $k \leftarrow \text{argmin}_i\, p_i$
4. try to purchase goods: $q \leftarrow \min(d(p_k), A/p_k)$
5. if successful, consume goods and receive reward: $R \leftarrow R + r(p_k, q)$
6. when next timestep, goto 1

Bottom Line: Buyers

Behaviour: after acquiring selection fo price quotes for a timestep,

- buyers maximize expected reward
- by selecting seller with lowest price known
- purchasing their full demand for that price, or as much as they can afford
- buyer's reward account is then updated

Model Rationale

Buyers: two types of informational preferences

1. one buyer collects all the supermarkets' leaflets
2. the other buyer picks up the first one found
3. performing the purchase may be costly (e.g. drive t supermarket)
4. greedy buyers; short-term optimization

Constant Production, Constant Price

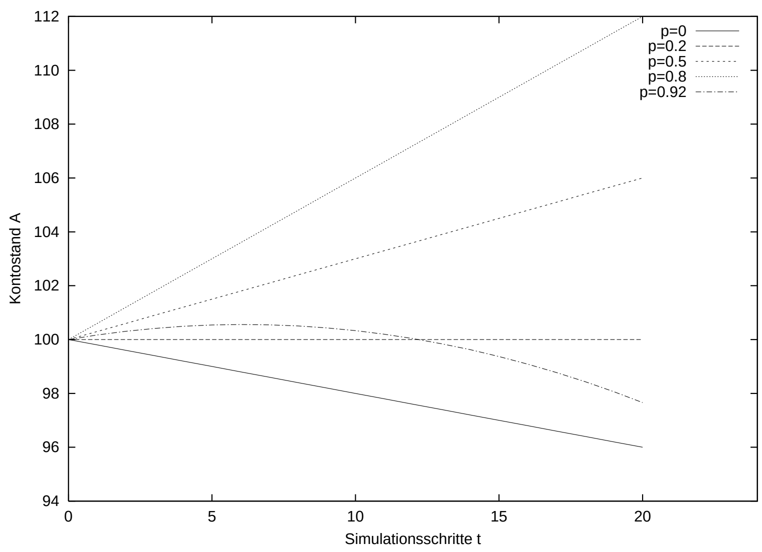

In the Following

Assume:

- number of buyers $M = 20$
- buyer income $i = 0.25$
- buyer maximum demand $d_{\text{max}} = 1.0$
- buyer maximum bid level $p_{\text{max}} = 1.0$
- seller initial credit $c = 100$
- seller initial price $p = 0.2$

Analysis: Market with Fixed Production

Model:

- production level P fixed
- only control of price p
- inflexible production chain, only price can be changed
- price adaptation by observation of inventory at each time
- fixed storage costs per unit per time step: tradeoff T sought between covering store costs and losses due to unability of serving all buyers' orders

Heuristics for T

Model: lower prices assumed to lead to increased sales

- raise price whenever $S < T$
- lower price whenever $S > T$

Here:

- number of sellers $N = 4$
- production level $P = 1$
- inventory threshold $T = 2$
- price adaptation step $\Delta p = 0.01$

Effects

Effects: increasing transparency ω increases competition, tightens up price trajectories

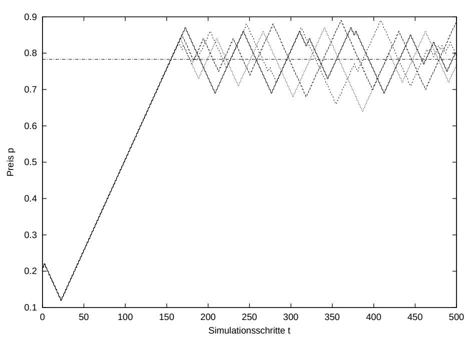

$\omega = 0$

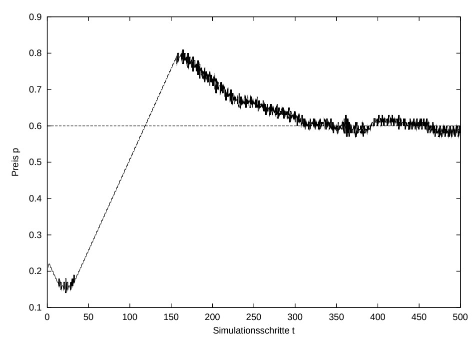

$\omega = 1$

Clearing Price

Model:

- individual production fixed, therefore aggregate production can be added up over sellers
- if sellers would sell all this quantity, total supply is same as total production
- one has for total production q^* and market clearing price p^*:

$$q^* = NP = 4.0 \qquad p^* = 1 - \frac{q^*}{M} = 0.8$$

- for low competition ($\omega \leq 0.3$), p^* quickly and independently reached by all the sell
- for higher competition, price is significantly lower than market clearing price

Collective Phenomenon!

Mechanism:

- concentration of market demand to cheapest seller: its inability to serve high demand due to production limitations
- excess supply of other sellers is taken into next time step and exerts pressure on the price
- market clearing price is reached quickly, but price decays later on:
  1. early on, sellers share demand equally
  2. as market clearing price is reached, fragile balance breaks up
  3. price wars and excess supply, pressing prices

Side Notes

Note: no modelling of opponents or even short-term profits

For Sellers:

- Prisoner's dilemma
- Nash equilibrium, not Pareto-optimal
- only reachable by collusion (cartel), illegal in typica market societies

Variable Supply

Model: control of production to satisfy all demands; control of profit via prices

Derivative Follower (DF): observes profit change depending on last price change to adapt prices

Observation: several DFs reach common price near monopolistic price, as long as no conflicting pricing strategy is introduced into market

Myopically Optimal (MO): perfect knowledge of buer demand and competitor price settings; use this information to select price yielding short-term maximum profit

Observation: in competitive scenarios with discrete fixed price levels, this leads to undercutting and cyclical price wars from monopolistic price to competitive equilibrium; far superior to DF

Price-Profit Learners

Price-Profit (PP): adaptation mechanism based on single-state Q-learning

- observed profits used to adapt profit expectations f current price
- choose new price through stochastic selection mechanism based on these expectations
- no information about composition of buyer population or competitor prices (as MO)

Idea: demand and competitor's pricing should be reflected in the PP relation

Observation: undercutting, outperforming DF; with less information, though, inferior to MO

Price-Profit Learners Competing

Question: how would strategies co-develop? Multiagent reinforcement learning

Observation: slower learning for convergence; resulting curve lies between monopolistic and Bertrand equilibrium price, depending on competition

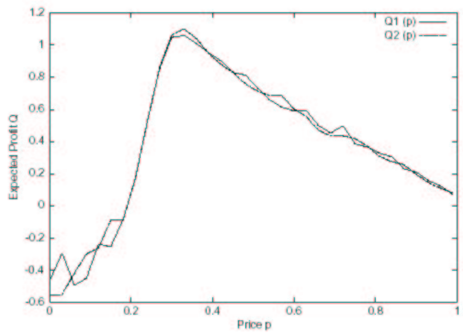

RL Pricing with Competitor Information

Model: Q-learning; state is given by lowest price on the market

Note: this is a full description of the learning task; only own and lowest competitors' price determine the demand

Expect: RL sellers to react more rapidly to situation than P

Experiments: no lookahead for Q-learning

Results: undercutting best-reply against any of the fixed strategies; for $\gamma = 0$ becomes very similar to MO

Price War Trajectories ("Nashian" Dynamics)