AUTOMATED MECHANISM DESIGN

Overview

- Auctions are increasingly used to buy and sell goods.
- Experience shows that “one size fits all” is not true.
  - Australia & New Zealand Cable TV auctions
- Hard to identify all the ways the auction might be gamed.
  - California power auctions.
  - FCC spectrum auctions.
- Designing auctions is a tricky business.

Mechanism design

- Mechanism design is the economics of designing auctions (and other games).
- Standard approach is to:
  - establish the optimum strategy of each player,
  - calculate the resulting equilibrium; and
  - check this meets the design criteria for the auction.
- Typical to assume every player plays her equilibrium strategy.
- From this *trading behaviour*, establish the best mechanism analytically.

Problems with mechanism design I

- The standard approach revolves around the concept of *Nash equilibrium*
- For a two player game a strategy (i^*, j^*) is a *Nash equilibrium solution* to the game (A, B) if:

$$
\begin{aligned}
\forall i, a_{i^*,j^*} &\geq a_{i,j^*} \\
\forall j, b_{i^*,j^*} &\geq b_{i^*,j}
\end{aligned}
$$

- For A to compute this, it needs to know the payoff to B.
  - Typically it won't.
- If A only knows the *type* of B probabilistically, we can use the notions of *Bayesian* Nash equilibrium.

Problems with mechanism design II

- The (Bayesian) Nash equilibrium cannot be computed for some (interesting) games.
  - Double auction.
- In some cases the BNE can't even be learnt.
- Even when the equilibrium can be computed, the players might not play to it.
- Maybe alternative notions of equilibrium are necessary?
- These thoughts have led to *computational economics*.

Problems with mechanism design III

- Assumes that players' best response respects the mechanism.
  - Nobody tries to game the system.
- Real life tells us that this is a BAD assumption.
- Maybe we can evolve mechanisms that are *strategy proof*.
  - Let lots of different strategies try out against our auctioneer.
  - Let the auctioneer *co-evolve*.
  - The final auctioneer should be able to cope with all possible gaming techniques.

WHAT WE DID (1)

Scenario

- Discriminatory price double auction simulating electricity market.
- NB buyers and NS sellers.
- Each B_i and seller S_j has a *generating capacity* GC_{B_i} or GC_{S_j}.
- Market parameters:

$$\text{RCON} = \frac{NS}{NB}$$

and

$$\text{RCAP} = \frac{\sum_{i=1}^{NB} GC_{B_i}}{\sum_{j=1}^{NS} GC_{S_j}}$$

define the relative balance of power in the electricity market.

Scenario 2

- Buyers and sellers try to maximise their local profit:

$$Profit_{B_i} = \sum_{k=1}^{NT_{B_i}} private\_value_k - trade\_price_k$$

and

$$Profit_{S_j} = \sum_{p=1}^{NT_{S_j}} trade\_price_p - private\_value_p$$

- Auctioneer tries to maximise efficiency:

$$ME = \frac{GlobalProfit}{TP}$$

where

$$GlobalProfit = \sum_{i=1}^{NB} Profit_{B_i} + \sum_{j=1}^{NS} Profit_{S_j}$$

and *TP* is the profit when every trader bids at its private value.

Scenario 3

- In each round buyers bid p_{b_i}, sellers ask s_{a_j}.
- Auctioneer matches overlapping bids and asks.
- Pricing rule sets the transaction price in the interval:

$$[p_{a_j}, p_{b_i}]$$

- What should traders bid to optimise their local profit?
- What is the optimum pricing rule to maximise global profit and/or efficiency?
- Try to evolve these using genetic programming

What we did first

- Initally we calibrated against results by Leigh Tesfatsion.
- Then we coded up the buyers and sellers as GPs.
- The auctioneer used a simple discriminatory price rule.
- Evolve traders:
  - Start with random population;
  - Run many auctions;
  - Breed the best traders, judged on their average profit; and
  - Repeat.
- Can this generate high efficiency trading?

Results for standard auctioneer

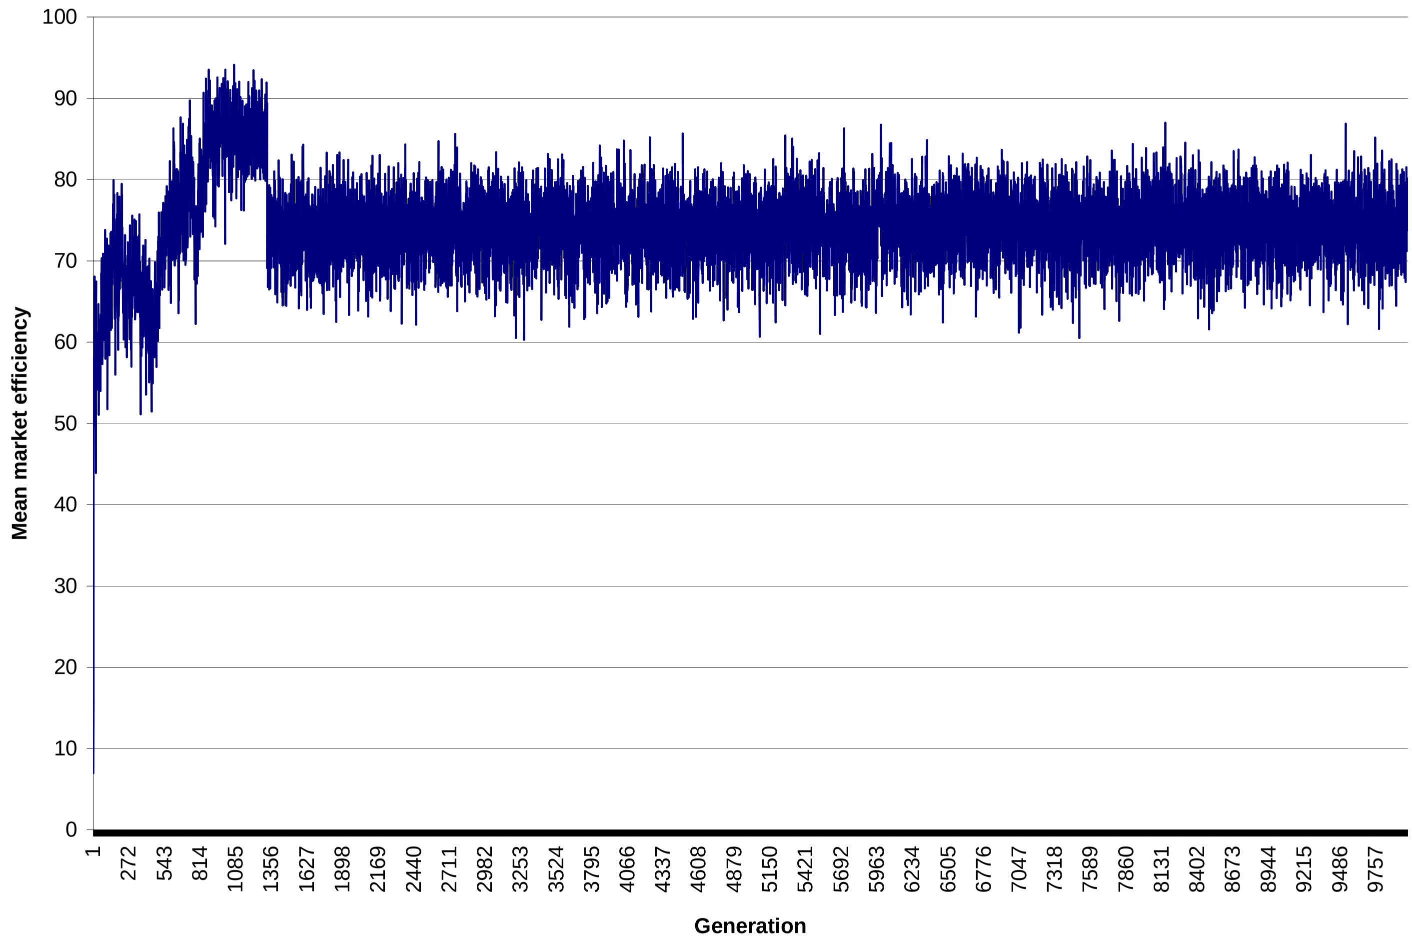

What we did second

- Now have the auctioneer as a GP as well.
- Do the same learning as before for the traders, but evolve the auctioneer at the same time.
- Evolve auctioneer:
  - Start with random population;
  - Run many auctions;
  - Breed the best auctioneer, judged on the global profit; and
  - Repeat.
- Can this generate high efficiency trading?

Results for evolved auctioneer

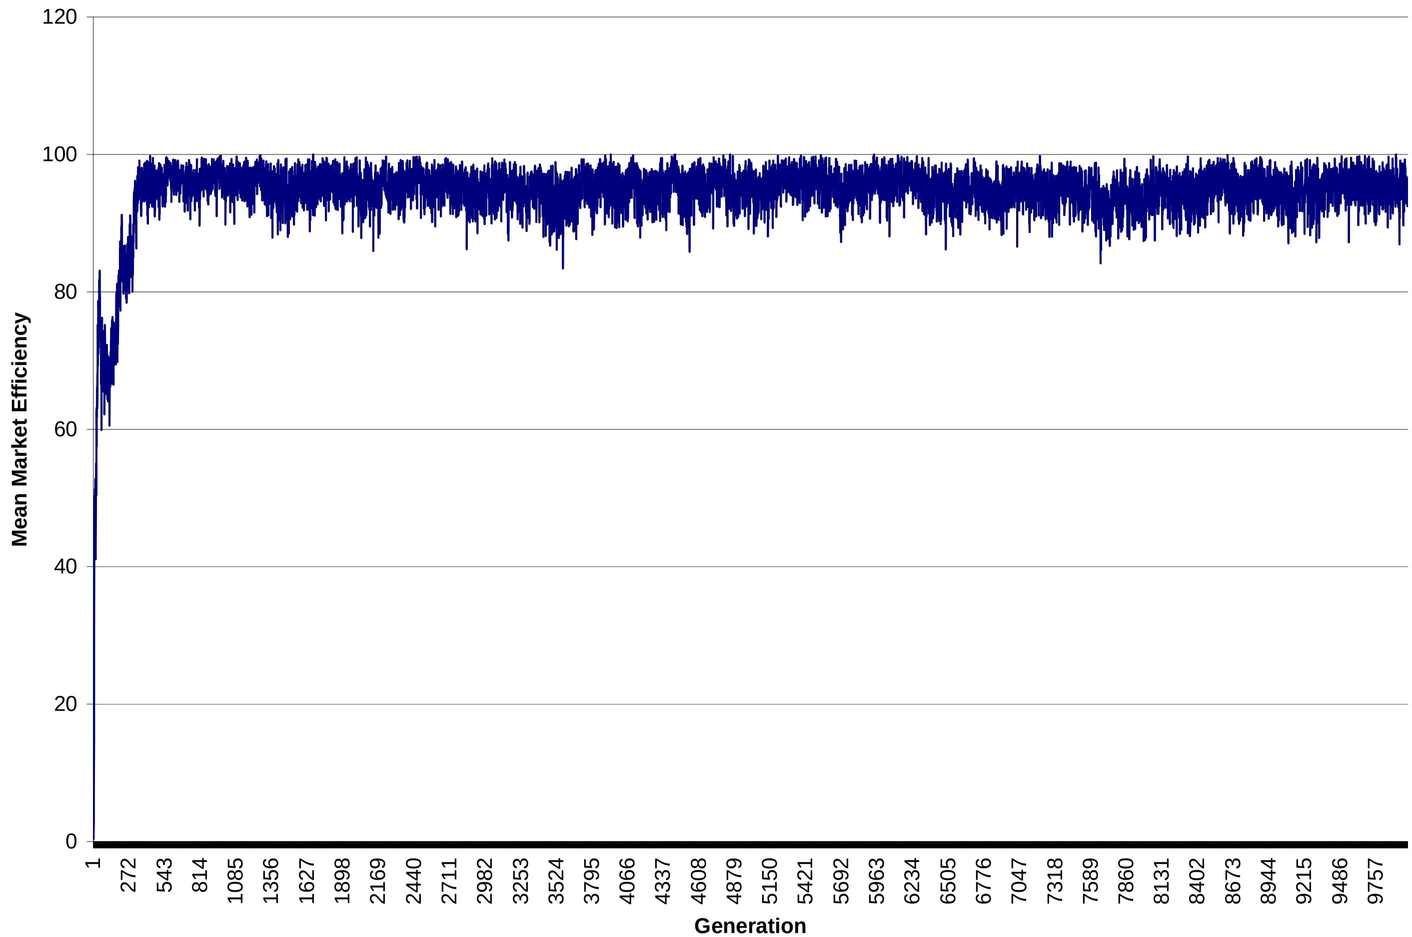

Problem

- Though we get high efficiency, this might not be for a good reason.
- Seems as if the auctioneer is learning the distribution of private values, not how to respond to bids.

WHAT WE DID (2)

Scenario

- Discriminatory price double auction.
- NB buyers and NS sellers each trade 10 units.
- In each round buyers bid p_{b_i}, sellers ask s_{a_j}.
- Auctioneer matches overlapping bids and asks.
- Pricing rule sets the transaction price in the interval:

$$[p_{a_j}, p_{b_i}]$$

- What is the optimum pricing rule?

Scenario II

- Measures used:
  - Efficiency:

$$EA = 100\left(\frac{PBA + PSA}{PBE + PSE}\right)$$

  - Buyer market-power:

$$MPB = \frac{PBA - PBE}{PBE}$$

  - Strategic buyer market power:

$$SMPB = \frac{PBA - PBT}{PBE}$$

- Rate auction using linear combinations:

$$V = \frac{\widehat{EA}}{2} + \frac{\widehat{SMPB} + \widehat{SMPS}}{4}$$

What we did fourth

- What is the fitness landscape?
  - Take the standard pricing rule:

$$p_t = kp_{a_j} + (1-k)p_{b_i}$$

  - Work out rating for all values of k as buyers and sellers learn.
- Shows us the shape of the space in which we are trying to evolve the auctioneer.

Fitness landscape I

- Overall rating (30 buyers, 30 sellers)

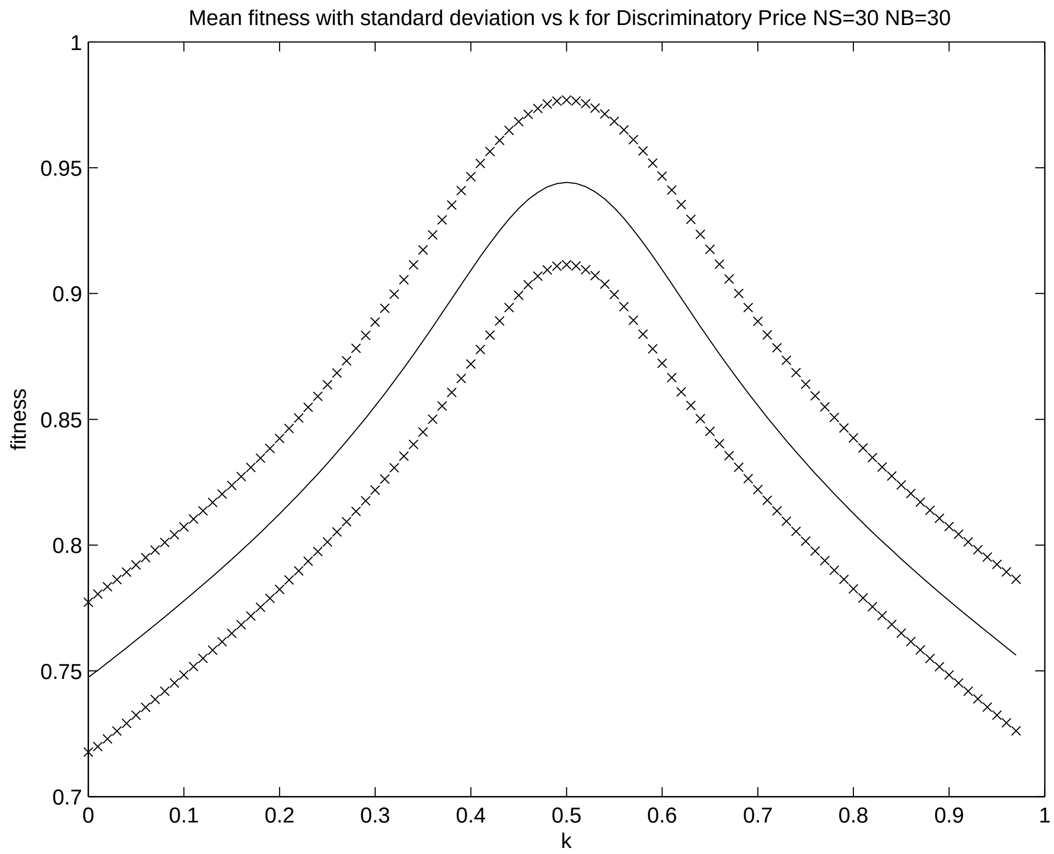

Fitness landscape II

- Overall rating (6 buyers, 6 sellers)

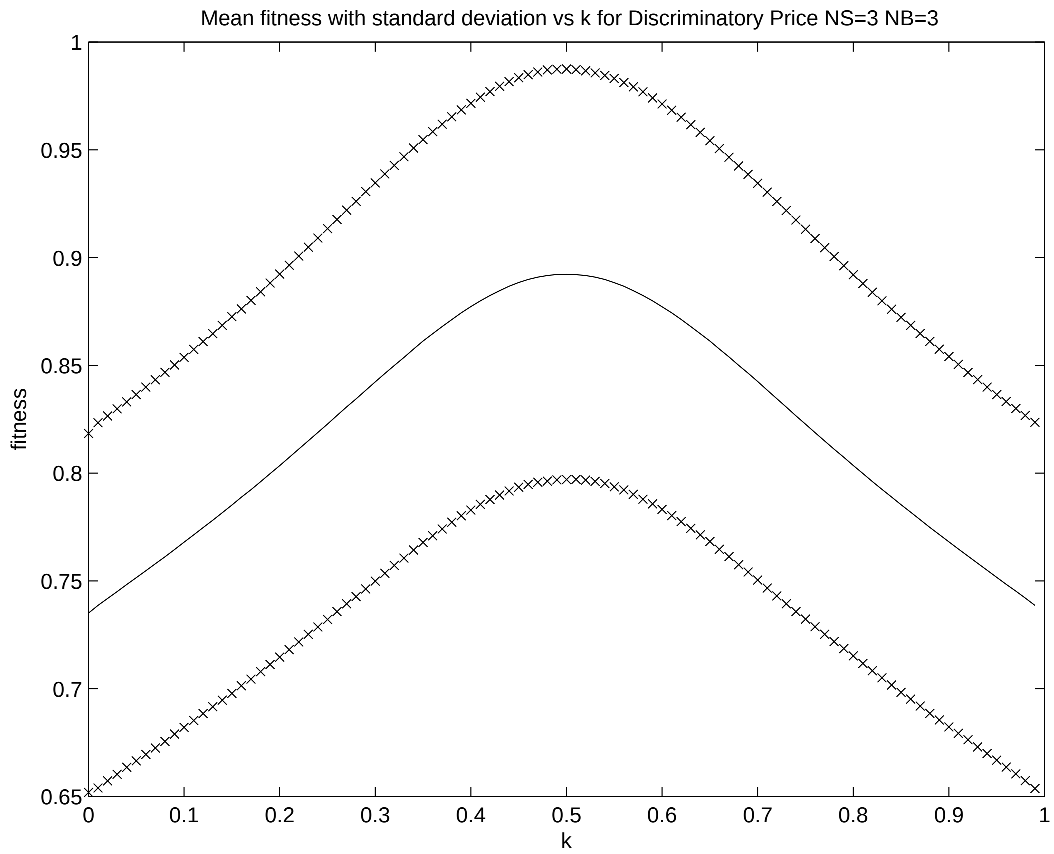

Fitness landscape III

- Strategic buyer market power (30 buyers, 30 sellers)

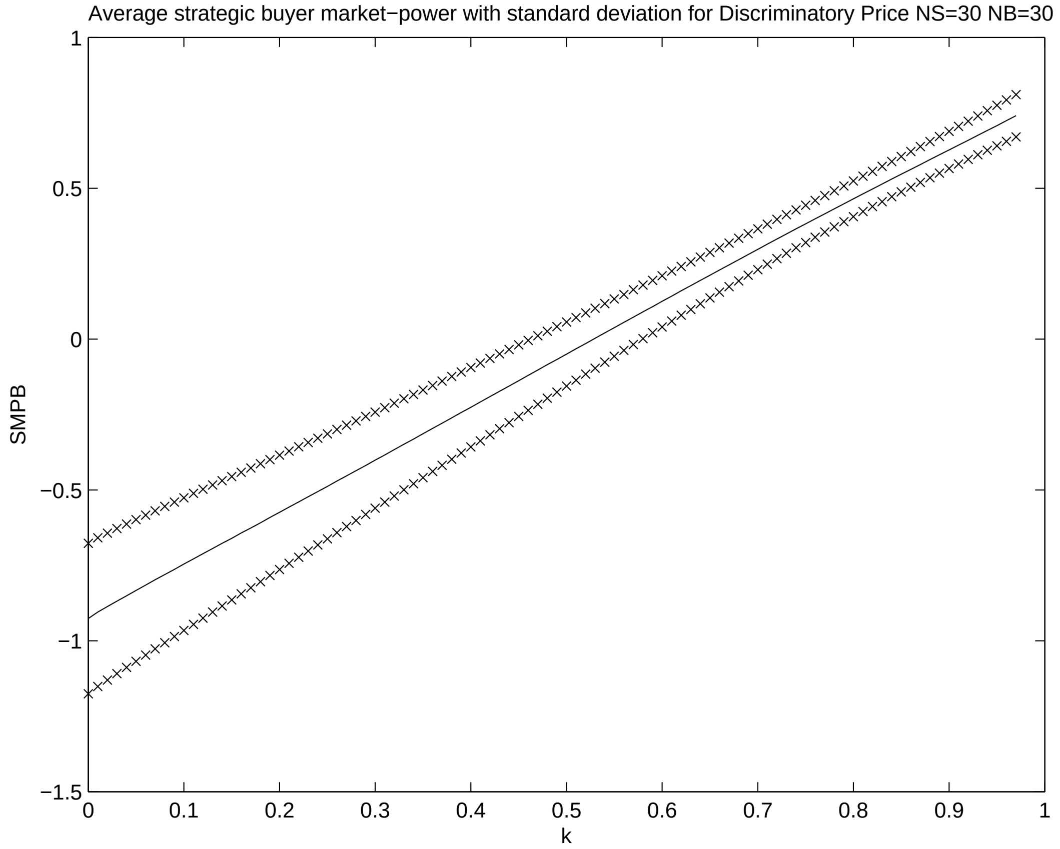

Fitness landscape IV

- Strategic seller market power (6 buyers, 6 sellers)

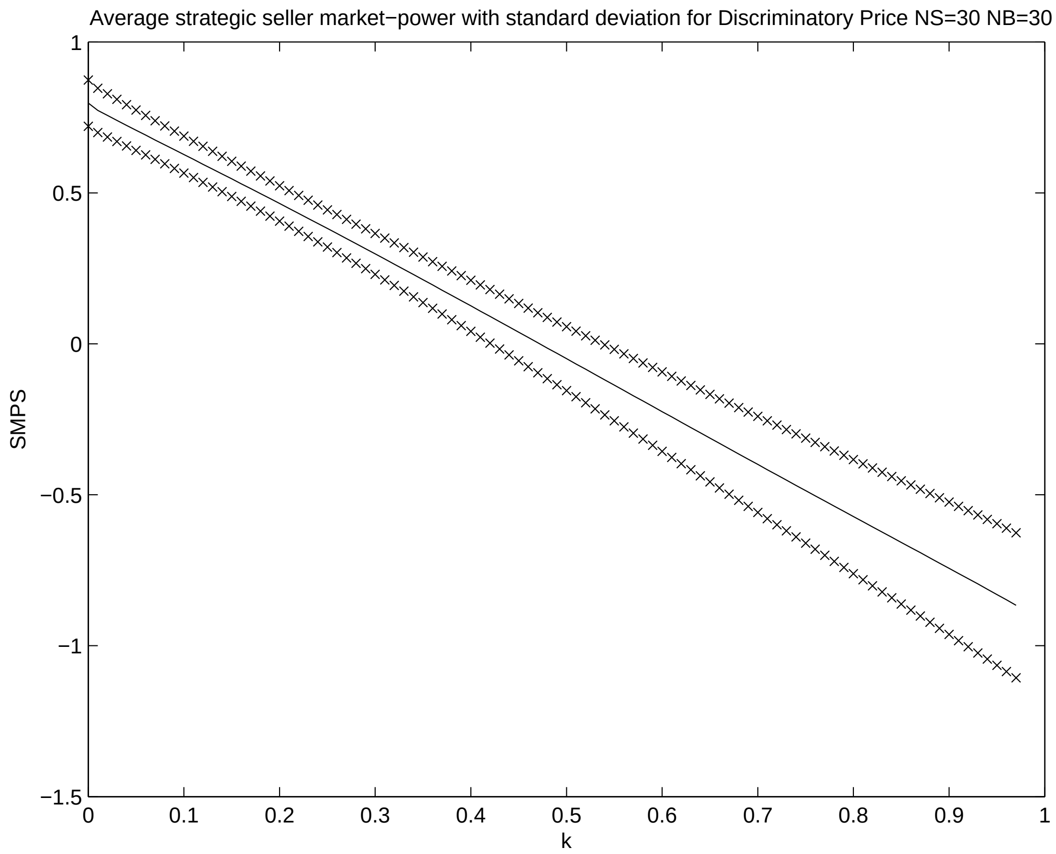

What we did fourth

- Take a learning model of trading behaviour.
- Take a mechanism for deciding prices based on offers.
- Let agents learn how to trade in this market.
- Evaluate the market against some criteria:
  - Efficiency
  - Pareto optimality
- Change the mechanism & repeat.

  In other words we learn the mechanism against a set of traders that themselves learn to play the mechanism.

Evolving a pricing rule

- After 90 generations

$$
\begin{gathered}
((0.6250385(0.93977016(\mathit{ASKPRICE} + 0.76238054)))+ \\
(((((-0.19079465)/(\mathit{ASKPRICE} - (((\mathit{BIDPRICE} \\
+\mathit{BIDPRICE})/(((((\mathit{ASKPRICE} - 1) + 1.6088724)/ \\
(((1 - \mathit{ASKPRICE}) - (\mathit{ASKPRICE}/\mathit{ASKPRICE}))+ \\
(2.5486426 + (\mathit{BIDPRICE} + 0.000012302072))))+ \\
((\mathit{BIDPRICE}/\mathit{ASKPRICE}) + ((\mathit{BIDPRICE} + \mathit{BIDPRICE}) \\
+(1.430315)/(\mathit{BIDPRICE} \cdot \mathit{ASKPRICE})))))\mathit{ASKPRICE})) \ldots
\end{gathered}
$$

Evolving a pricing rule II

- Price plotted against p_a and p_b.

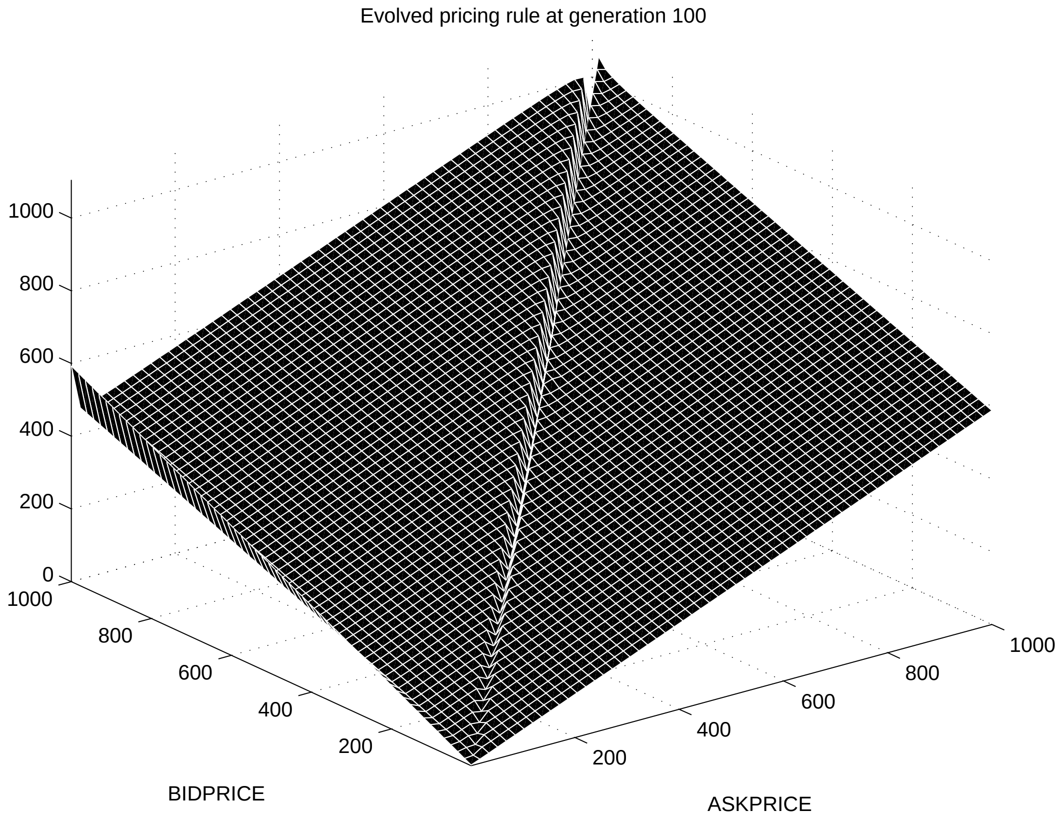

Evolving a pricing rule III

- Deviation from $k = 0.5$.

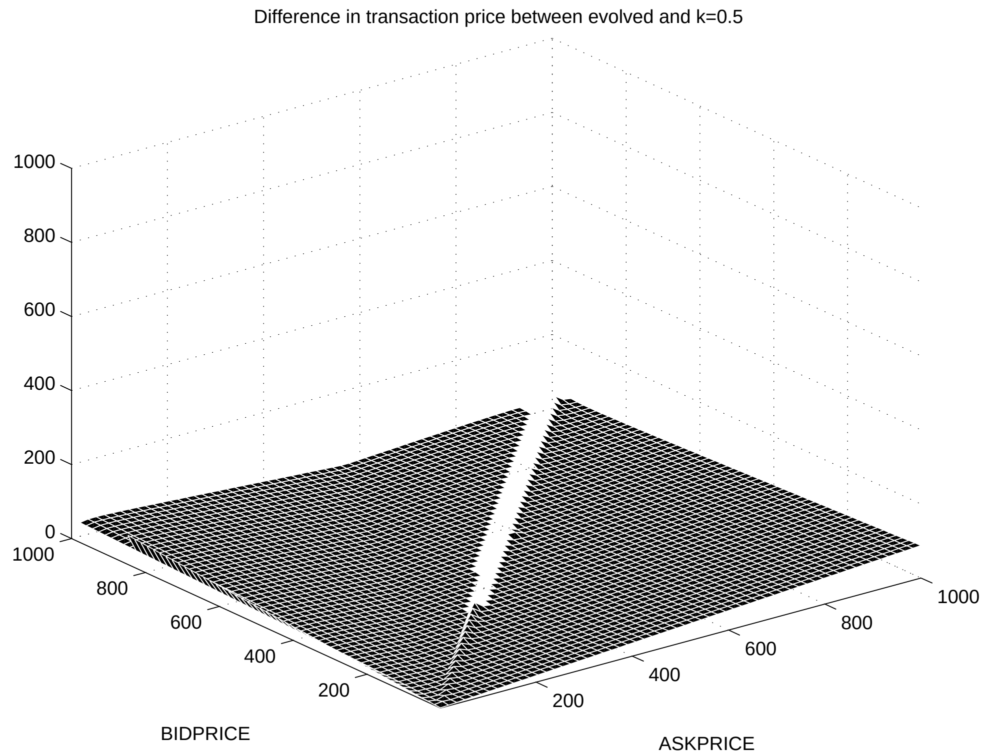

Summary

- First experiments seem to reinforce the idea that “the market is the thing”.
- Reasonable efficiency even given dumb bidders.
- However, can’t learn smart auctioneer from dumb bidders.
- Second experiments, suggest that balancing demands of buyers and sellers points towards $k = 0.5$ auction.
- With a nicely defined fitness landscape, we can recover a sensible pricing rule.