ESHIA-S4 Summer School, La Londe les Maures (September 2007)

Modelling learning on 'simple' markets

Eric DARMON

eric.darmon@univ-rennes1.fr

CREM-CNRS and Université de Rennes 1

UNIVERSITÉ DE RENNES 1

Agregation on markets topics…

Several talks on AB Markets by

A. Kirman
+ C. Deissenberg
+ J. Rouchier
+ J. Arifovic
+ this talk

Focus on technical and practical aspects of AB market simulations

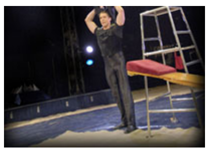

Back to Kirman's talk

- There is no reason to believe that markets should be 'spontaneously' at equilibrium

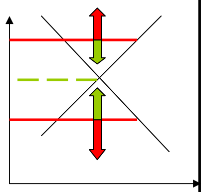

- So…how do markets « adjust » ?
  - Standard view
    - Processus of tatonnement
    - Processus of non tatonnement : better ?

  - Decentralized Pairwise meetings
    - Posted versus secret/bargained prices: Brenner 02'
    - Repeated versus one-shot purchases
    - → **Theories of Market organizations**

Which market ?

- 2 ways of considering things
  - Make it as « realistic » as possible
    - May seem more satisfactorily
    - But raises, many methodological issues:
      - What causes what ?
      - How to calibrate the model ?
    - Can be used only when compared to actual/observed data
  - Application of the KISS principle
    - **K**eep **I**t as **S**imple as **P**ossible
    - Leaving aside many features
    - Computational approach

Which market …

- For what kind of result ?

1. **Comparing two outputs of the same agent-based model**
   - Cf. talk of Y.Arifovic: comparison of two variants with different information sets on a call market
   - Another illustration : a basic commodity market with two types of interaction processes: **endogenous** versus **random** matching of buyers/sellers

Anonymous versus Non anonymous markets

Endogenous matching

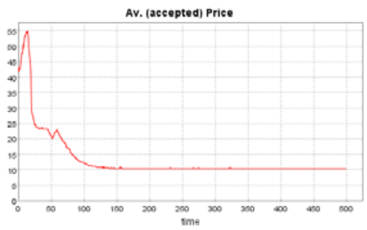

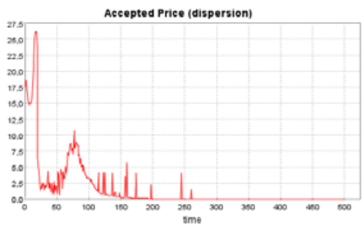

Anonymous matching

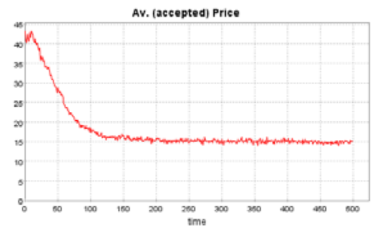

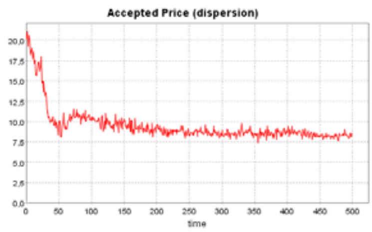

Which market …

- For what kind of result ?
  1. Comparing two outputs of an agent model
     - Cf. Y.Arifovic: comparison of two variants with different information sets
     - Another example: a basic commodity market with two types of interaction processes: **endogenous** versus **random** matching of buyers/sellers
  2. Comparing the output of an agent-based model to that of an already known model
     - Economic litterature: mainly relies upon the Nash Equilibrium concept to deal with interactions
     - Yet,
       - Nash may be complicated to compute for simple agents.
       - Agents may not have a complete knowledge
     - Convergence or divergence ?

A simple *Search* Market

- One particular market story …
  - Non-frequent purchases
  - Posted prices: information on posted prices is not publicly available
    - The number of buyers/sellers interactions is not limited yet buyers have to pay a cost (travel, opportunity cost) to get it
- Game theoretical analysis
  - Search theory : Stigler [1969]
  - Sellers **set** prices: price makers yet contrained by the behavior of other sellers

A simple Search Market

- **Key points**
  - When information is costly…
    - Buyers: may stop searching before visiting the whole set of sellers: high search cost buyers will visit less than lower ones: heterogeneous behaviors
    - Sellers: anticipate buyers' behavior : know that they can propose different prices to each types of buyers
  - From that, sellers can infer the equilibrium price distribution and propose …
    - Pure price strategies (Salop & Stiglitz, 1977) : one seller posts one price only
    - Mixed price strategies (Varian, 1980) : one seller samples price from the equilibrium price distribution
- **Main Conclusions**: no convergence to some 'competitive' price, price dispersion for homogeneous items

A simple Search Market

- **What's the problem ?**
  - Decentralized setting …
  - Yet just apparently : what do Sellers need to know to compute their strategies ?
- **Common feature of these models** : static notion of equilibrium where sellers have a perfect knowledge (common knowledge hypothesis)
  - About buyers' characteristics
  - About sellers' ones
- … but what if these features are not known ?

A simple Search Market

- Comparing Agent-based to equilibrium models:
  - **Why ?**
    - **Relevance of the Nash Search Equilibrium (NSE) in mixed strategies:**
      - NSE is too complicated to be played by 'real' sellers. Yet, can it be learned ?
    - **IO/Competition policy:** is Nash is never reached, can we predict seller price reactions w.r.t. more and better informed consumers?
  - **How ?**
    - Do we (or can we ?) compare the same items ?
    - Can we generalize ?

Outline of the talk

- The model**s** : presentation & implementation
- The comparison in the reference case :
  - Main results
  - Behind the comparison: some inherent difficulties and limits
- Generalization & Validation issues
  - Robustness towards the learning algoritm
  - Robustness towards Buyers' search process
- Next …
- Related topics

The Model – Simple posted price market

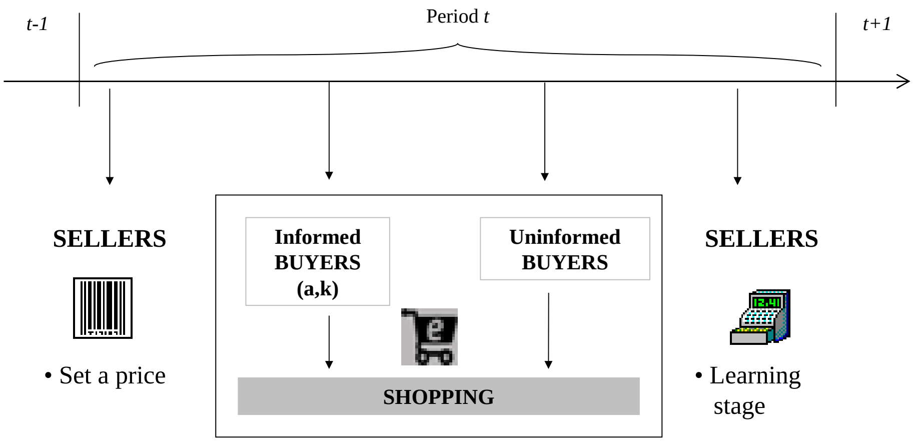

The Model – Buyers

- Need to buy one unit at each period at a maximum reservation price (***v = 100***)
- Non-frequent purchases
- Two types of buyers :
  - **Uninformed Buyers** : visit randomly one seller and shop if the price proposed is less than or equal to *v*
  - **Informed Buyers** : sample randomly **k** sellers, and shop at the best price (if less than or equal to *v*)

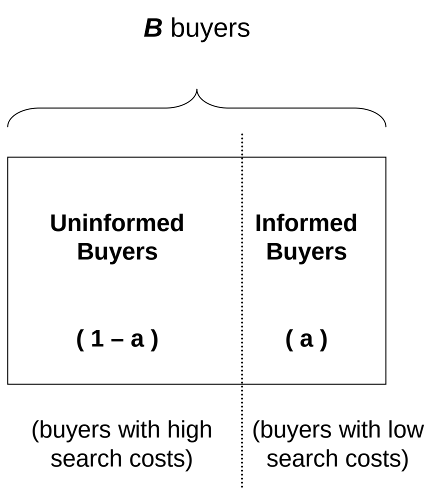

Simplest version: fixed sample size rule hypothesis

The Model – Sellers (I)

- Fixed population of ***S=20*** sellers
- One decision per period : price (p_t)
  - No bargaining; fixed price within a time period
  - Simultaneous choice with other sellers
  - Sellers know that the maximum willingness to pay for the good is *v* , hence $p_t <= v$
  - Homogenous production cost *c* (hence *c* = 0)

v = 100
99
98
…
***P*$_t$?**
…
2
1
c = 0

How to set price ?

v = 100
99
98
…
***P*$_t$?**
…
2
1
c = 0

The Model – Sellers (II) : How to set price ?

- **Equilibrium model**: sellers know *a* and *k*, and maximize their expected profit conditional on the profit maximization of other sellers
  - Nash equilibrium

- Additional hypothesis: symmetric equilibrium posted price distribution

- **Results** (proofs in Waldeck, JEBO *2006*) : a bimodal continuous price distribution

Sellers' strategies at equilibrium

- Bimodal distribution: two main strategies
  - High price/low sales
  - Low price/high sales

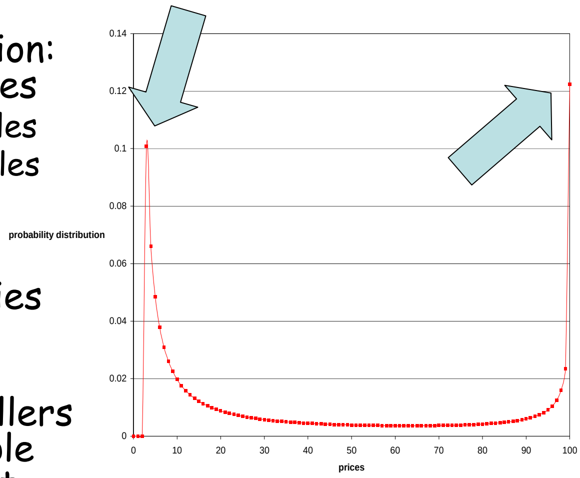

Chart : NSE distribution of posted prices with a=0.9 and k=5

- All these strategies lead to the same expected profit
- To set a price, sellers just need to sample randomly into that distribution

How to set price ? [Cont'd]

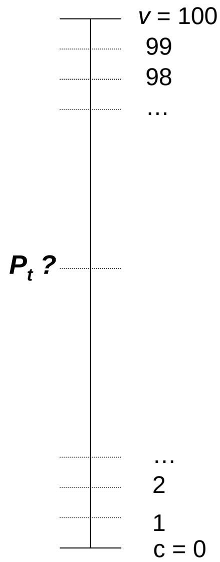

The Model – Sellers (II) : How to set price ?

- **Equilibrium model**: sellers know *a* and *k*, and maximize their expected profit conditional on the profit maximization of other sellers
  - **Results:** a bimodal continuous price distribution

- ***Agent-based model: which theory of learning ? [cf. C.Deissenberg & J.Rouchier]***
  - Fictious play with bayesian updating ?
  - Simple (naïve) imitation rule ?
  - Reinforcement learning ?
    - Sellers do not know much about the characteristics of their environment
    - They use their own experiences to set « better » prices

The Model – Sellers (II) :
How to set price ?

- Reinforcement Learning (Sutton, 98) :
  - **General principle:** sellers tend to use the price strategies that exhibited the largest payoffs in the past
  - **3 key ingredients**
    1. A set of available strategies and a reward (profits)
    2. A method to initialize fitness at the first step and another to update them when the strategy is selected
    3. A method to select the strategies used

Which application on a search market ?

RL: A set of available strategies (I)

Condition	Rules (R_i)	Fitness (F_i)
	rule #100 : Set price 100	F_{100}
	rule #99 : Set price 99	F_{99}
	rule #98 : Set price 98	F_{98}
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	rule #2 : set price 2	F_2
	rule #1 : set price 1	F_1
	rule #0 : set price 0	F_0

More sophisticated strategies could be used:

Example: two or multi-periods strategies (on conditions and/or actions)

RL: A set of available strategies (I)

Condition	Rules (R_i)	Fitness (F_i)
	rule #100 : Set price 100	F_{100}
	rule #99 : Set price 99	F_{99}
	rule #98 : Set price 98	F_{98}
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	...	**...**
	rule #2 : set price 2	F_2
	rule #1 : set price 1	F_1
	rule #0 : set price 0	F_0

• Criterion used to reward the rule currently used :

actual profit generated by that rule

$$\pi_t = n_t \left(p_t - c \right)$$

RL: Initialization & Update rules (II)

- How are fitness initialized ?

$$F_{t=0}^{i} = \delta\left(Bv\right) \text{ with } \delta \in \left]0,1\right]$$

- How are they updated ?

$$F_{t+1}^{i} = (1-\alpha)F_{t}^{i} + \alpha\pi_{t}^{i} \text{ with } \alpha \in \left]0,1\right]$$

Weight attached to the latest experience

Current profit

RL: Selection rule (III)

- How are price strategies selected ?
  - Simplest way: greedy or epsilon-greedy selection mode

    Not efficient for a high number of rules
  - Exploration-exploitation tradeoff
  - Stochastic selection mode: Boltzmann i.e. logistic distribution)

$$prob\{\text{ Set Price } i\} = \frac{e^{\frac{F_t^i}{\tau}}}{\sum_j e^{\frac{F_t^j}{\tau}}}$$

  - **For the selection process (only)**, fitnesses are rescaled between 0 and 1

$$\text{with } \tau \in \left]0,+\infty\right[$$

Rescaling fitnesses: why ?

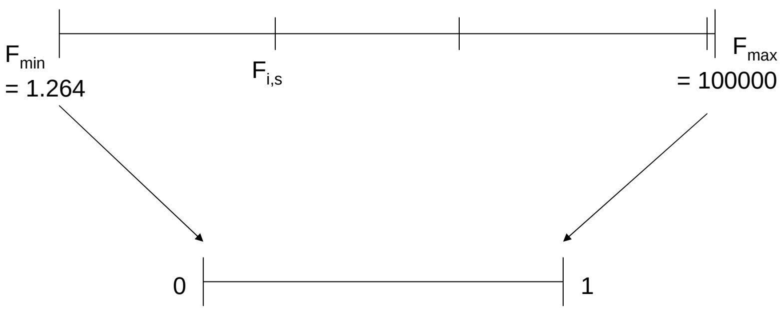

- Computational problems with exponential values
- ‘Temperature’ parameter independant of the absolute magnitude of the payoffs: making simulations comparable

Implementing the model

Java implementation: object-oriented programming: from agents to classes (attributes and methods)

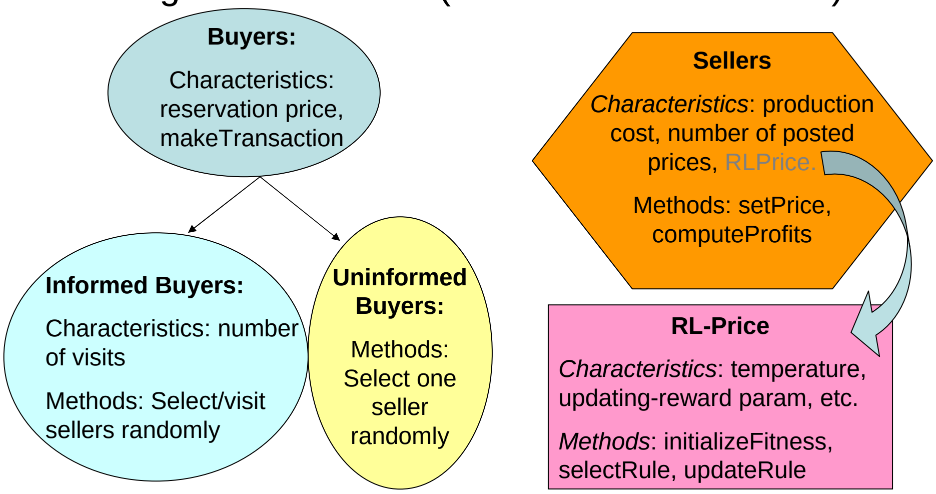

Class MarketDynamics:
environnement & scheduler

Attributes: all buyers and sellers
Methods: aDayInTheMarket, allDays

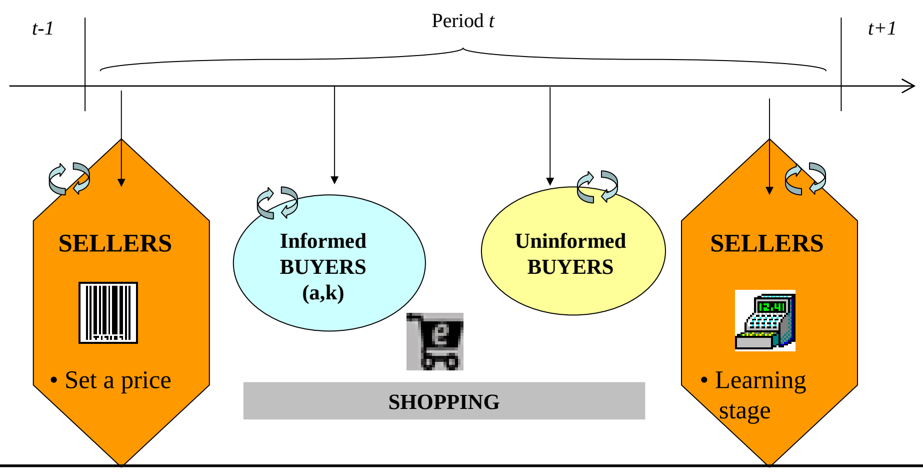

List of parameters & Technical features

Parameter	Default value
Number of sellers	20
Number of buyers	1000
Number of rounds	1000
a	variable
k	variable
Exploration-exploitation parameter	0.1
Fitness initialization	0.2
Reward-updating parameter	0.8

Buyers' search features: *a*, *k*

Learning parameters: Exploration-exploitation parameter, Fitness initialization, Reward-updating parameter

Java Program and pseudo code : http://e.darmon.free.fr/fssmarket/

Outline of the talk

- The models : presentation & implementation
- **The comparison in the reference case :**
  - **Main results**
  - **Behind the comparison: some inherent difficulties and limits**
- Generalization & Validation issues
  - Robustness towards the learning algoritm
  - Robustness towards Buyers' search process
- Next …
- Related topics

Let's « play » with parameters

- A first case

A first comparison of the two outcomes

- *At equilbrium →*
- *Comparison between this distribution and the* ***stationary*** *price distribution*
  - *We skip price dynamics voluntarily*

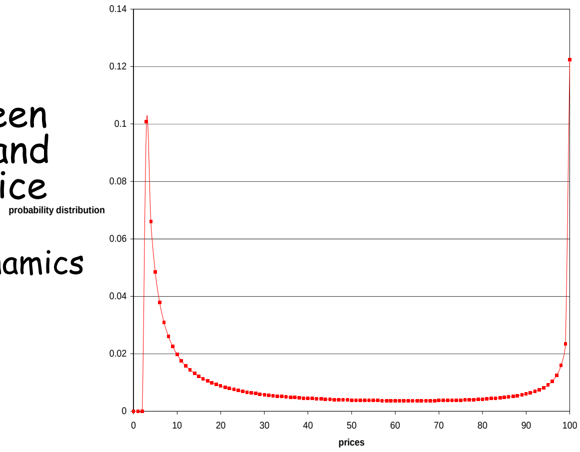

Chart : NSE distribution of posted prices with a=0.9 and k=5

Results (I) – Price distribution

- The stationary distribution with reinforcement sellers does **not** converge to the Nash distribution for parameters "a" different from 1 (for a=1, convergence depends on the time length of simulation).

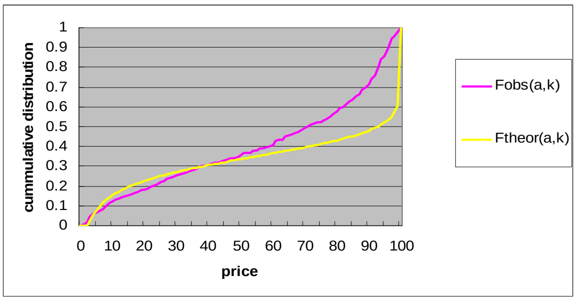

Chart : NSE and Learned cumm. distribution of posted prices with a=0.8 and k=10, 50 periods and 20 sellers

Let's « play » with parameters

- A first case
  - What if buyers are more informed ?
- What about learning parameters ?
  - Increasing memory
  - Increasing « temperature »
  - Magnitude of the initial fitnesses
- How to go beyond these intuitions ?

Sensitivity analysis and « multi-shot » simulations

Behind the previous comparison (I)

- Do we (can we) compare the same models ? Two inherent and critical issues
  - Continuous versus Discrete :
    - Nash: Price strategy as a continuous distribution
    - AB Model: need to discretize price to have a finite strategy set: what should be here the **'optimal' number of price strategies ?** 2 ? 3? 10 ? 100 ? 1000 ? « Infinity » ?
  - Strategy selection process
    - Introduces a noise into the simulated distributions
    - Which alternative ?
      - Decreasing arbitrarily the temperature coefficient ?
      - Refining the equilibrium concept: Quantal Response Equilibria instead?

Behind the previous comparison (II)

- An alternative to 'overcome' the non convergence results :
  - Open question: are such analytical models useful to predict actual behaviors on markets ?
    - On 'real' markets, many differences : in production costs, in the supplied quantities, in willingness-to-pay, etc.
    - Such equilibrium models are simply not made to describe how sellers would **actually** perform but rather to guess how they may **react to changes in market conditions** e.g. change in buyers' behaviors

→ Redefinition of the comparison criterion

Behind the previous comparison (III)

- Redefinition of the comparison criterion :
  - **Softer criterion:**
    - Reactions of adaptive sellers to the main structural parameters
    - Test whether these reactions are compatible with Nash predictions
  - **2 important parameters** here : proportion of informed buyers (**a**) and number of sellers visited by informed buyers (**k**)
    - What happens for the price dispersion and the average price if the proportion of informed buyers increases ?
  - Switch from quantitative to qualitative results

Nash predictions (II)

Change in buyers' search behaviour	Effect on Mean Price	Effect on Price Variance
• Increase in ***a*** (more informed buyers)	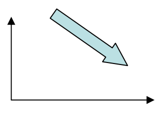	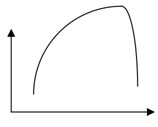
• Increase in ***k*** (informed buyers make more visits)	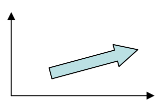	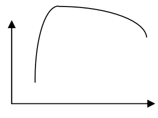

Sensitivity analysis

- From « one shot » to « multiple » simulations
  - Objective: test wether a simulation is « representative » of something
  - Example on Parameters *a* and *k*

Randomly chosen between 0 and 1

Randomly chosen between 1 and the number of sellers

  - For each simulation, we track « final » data: average price, price dispersion

Results (VI) – Price dispersion & coefficient *k*

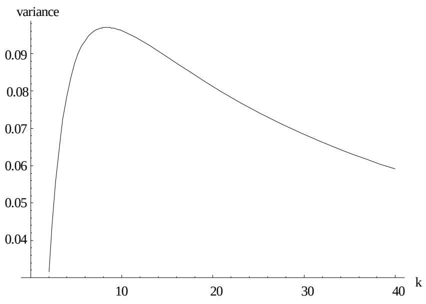

Figure 2: Variance of F(p) as a function of k (for a =0.5, v=1).

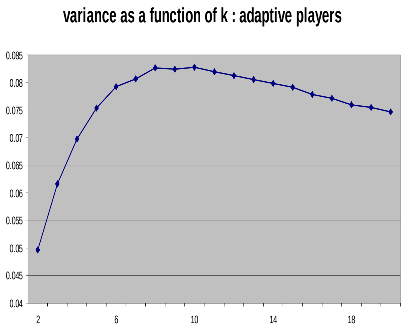

Conclusion: similar properties : inverse U- shaped curve, maximum price dispersion at around the same value of k!

Rk : except for a = 0.1 (where dispersion is constant)

Results (II) – Av. price & coefficient *a*

- Variations of posted price with respect to a : true for all ***a*** (a ≠ 1 and a ≠ 0)

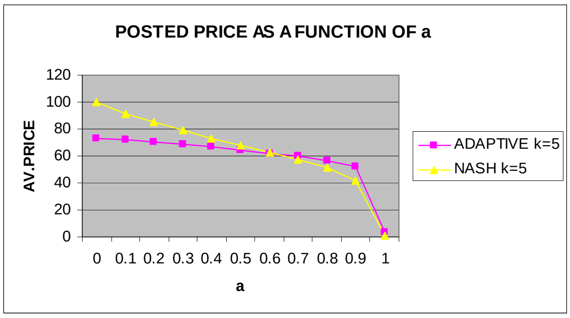

Chart : NSE and Learned distribution of posted prices,
average over 100 runs for each parameter configuration (a,k)

Results (III) – Av. price & coefficient *k*

- Variations of posted price with respect to ***k*** : True all k ≠ 2 (a ≠ 1 and a ≠ 0)
- For a = 1, the mean price decreases with k (mainly due to convergence time)

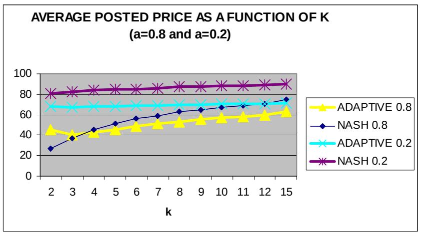

Chart : NSE and Learned distribution of posted prices,
average over 100 runs for each parameter configuration (a,k)

Results (IV) – Price Dispersion & coefficient *a*

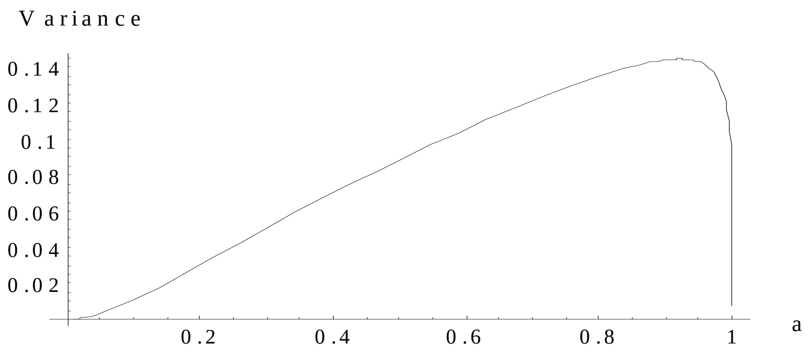

Figure 5: Variance of the price distribution as a function of a (for k=5, v =1)

Results (V) – Price disp. & coefficient *a*

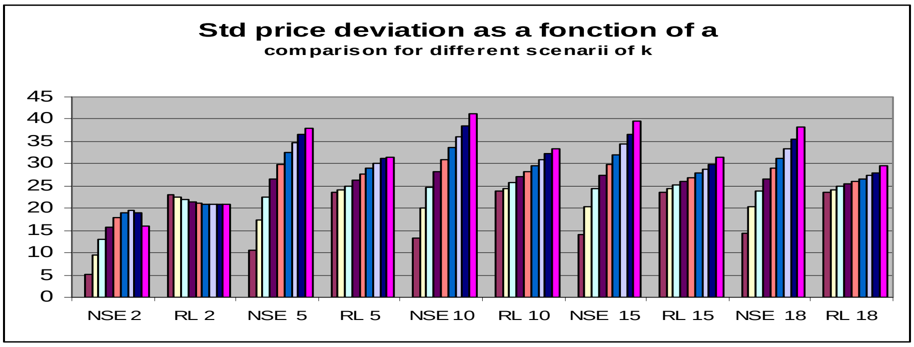

Conclusion : same skew to the right ⇒ price dispersion is in general increasing with a (true all k except ≠ 2)

Summing up: Nash vs RL

- Learning the NSE Price distribution ? ✘
- Predictive capacity of NSE with respect to a change in buyers' search behaviour

	Av price	Price Dispersion
a	☑ (except k=2)	☑ (except k=2)
k	☑	☑

Discussion

- Is Nash equilibrium a good predictor ?
  - The contras … (cf. Posh 97)
  - … & The pros ! (good predictor with respect to changes in Buyers' search behaviour)
- Going further:
  - **Behavioral Game Theory / Behavioral Economics perspective**: linking together RL, Nash and experimental results: need for a triple comparison
  - **M2M perspective** : can we generalize these results ? What about robustness ?

Discussion (I): Behavioral Game theory / Behavioral Economics perspective

- Erev & Roth (AER, 1998): experimental results
  - Nash equilibrium prediction performs poorly for constant sum games with a unique mixed Nash equilibrium
  - A simple model of RL performs better than Nash for normal form games with a unique mixed strategy equilibrium
  - Result valid for both descriptive and predictive issues
- But …
  - Alternative models exist and may explain data as well e.g. Sarin & Vahid (Games & Econ. Behav., 97) : optimizing myopic players
  - Salmon (2001, Econometr.) : econometric methods may miss to identify correctly the learning rule that was used initially to generate the data

Outline of the talk

- The model**s** : presentation & implementation
- The comparison in the reference case :
  - Main results
  - Behind the comparison: some inherent difficulties and limits
- **Generalization & Validation issues**
  - **Robustness towards the learning algoritm**
  - **Robustness towards Buyers' search process**
- Next …
- Related topics

Generalization & Validation

3 issues :

1. Sensitivity towards RL specification and parameters
2. Sensitivity towards alternative types of learning
3. Sensitivity towards alternative buyers' behavior

A- Sensitivity to RL parameters & specifications

- Three main parameters governing the RL process :

Parameter	Impact
δ	
τ	
α	

A- Sensitivity to RL parameters & specifications

- Three main parameters governing the RL process :

Parameter	Impact
δ	strongly affect the **shape** of the dynamics but weakly affect the stationary state
τ	
α	

A- Sensitivity to RL parameters & specifications

- Three main parameters governing the RL process :

Parameter	Impact
δ	strongly affect the **shape** of the dynamics but weakly affect the stationary state
τ	For low values : add « noise » at stationary state For high values (>0.5): convergence to random choices
α	

A- Sensitivity to RL parameters & specifications

- Three main parameters governing the RL process :

Parameter	Impact
δ	strongly affect the **shape** of the dynamics but weakly affect the stationary state
τ	For low values : add « noise » at stationary state For high values (>0.5): convergence to random choices
α	Affect the convergence time (especially for low values of *k* and *a*) ; special case α = 1

A- Sensitivity to RL parameters & specifications

- Other (yet unexplored) formulation of the RL algorithm :
  - Choice of another selection rule :
    - Simple Linear selection rule
    - Greedy action

$$prob\left\{\text{ Set Price } i\right\} = \frac{F_t^i}{\sum_j F_t^j}$$

$$\left\{\text{ Set Price i}^*\right\} \text{if } \mathrm{F_t^{i*}} = \max_i \left\{\mathrm{F_t^i}\right\} \text{with probability } \left(1-\varepsilon\right)$$

$$\left\{\text{ Set Price i} \neq \text{i}^*\right\} \text{with probability } \left(\varepsilon\right)$$

Individual or Social Learning ?

- Vriend: points out the differences between the 2 (spite effect)
  - Does it stand here ?
  - How to add social learning in this system ?
- Introduction of a simple naive RL process
  - Sellers observe what other sellers did
    - Posted Price …
    - … levels of transaction also
  - With that information, they could infer the profit they would have done if they have been using this pricing rule
    - What could he do with this piece of information ???

B- Individual or Social Learning ?

- Modification of the learning algorithm to allow for a simple social learning: naive social learning
  - At each period, seller *S* can observe the payoff of sellers « -*S* », and modify his fitness accordingly
  - Additional Hypothesis: others' experiences are possibly « discounted » relative to seller *S*'s own experiences

For the rules played by other sellers

$$F^{j}_{s,t+1} = F^{j}_{t} + \alpha \left(\pi_{t,s'} - F^{j}_{s}\right)$$

B- Individual or Social Learning : results with Social Learning

- Are previous results sensitive to the type of learning? (case #1)

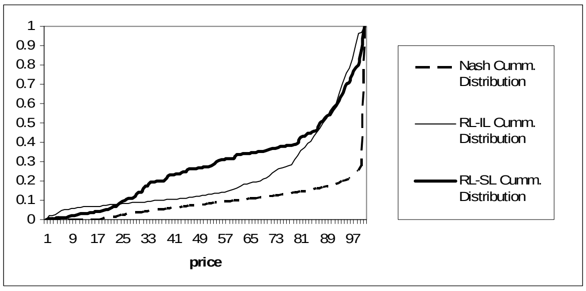

Main conclusions:

- Continuous learned price distribution with SL ('mixed strategies')
- Yet, no convergence to the IL case

Reported Case : $\{a = 0.2; \tau = 0.05; \delta = 0.2; \alpha = 0.8; \lambda = 1\}$

B- Individual or Social Learning : Results with Social Learning

- Are previous results sensitive to the type of learning ? (case #2) : a = 0.8 (more informed buyers) ; 2 different runs

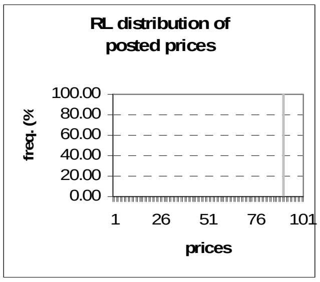

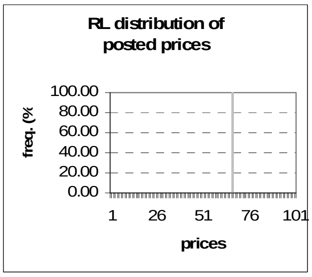

Main conclusions:

- Again, no convergence to the IL case
- Convergence to a single and path dependant outcome ('pure strategy')

Reported Case : $\{a = 0.8; \tau = 0.05; \delta = 0.2; \alpha = 0.8; \lambda = 1\}$

B- Summing up: Nash vs RL with SL

- Learning the NSE Price distribution ? ✘
- Predictive capacity of NSE with respect to a change in *a*/ *k*

Individual Learning

	Av price	Price Dispersion
a	☑	☑
k	☑	☑

Social Learning

	Av price	Price Dispersion
a	✘	☑
k	✘	✘

C- Generalization to other type of Buyers' behavior

- Buyers characterized by « Fixed Sample Size » search
  - What about **sequential search** ?
  - The number of visits is endogenous: buyers go on visiting new sellers as long as their expected benefit of doing so is positive
  - Again, two types of buyers

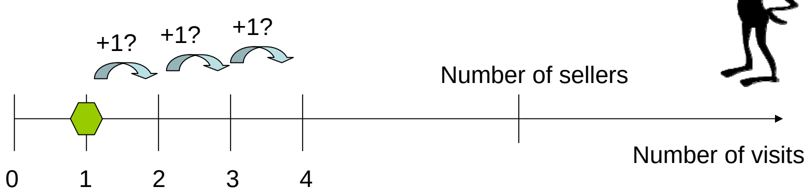

C- Generalization to other type of Buyers' behavior

- Sequential and FSS models
- Theoretical properties linking the two models
  - A famous result: at equilibrium, no search for high cost buyers in search models !
  - Do we retrieve these properties with adaptive agents ?

C- FSS or Sequential Search : results with Sequential Search

- Are previous results sensitive to the type of **search** ?

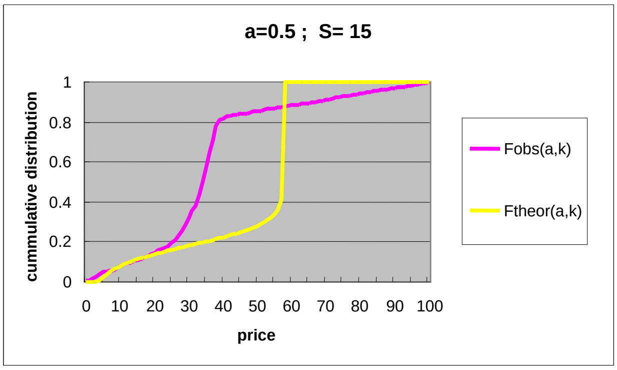

Main conclusion:

• No absolute convergence (as previously)

Summing up: Nash vs RL with SS

- Learning the NSE Price distribution ? ✗
- Similarity between Predictive capacity of NSE with respect to a change in buyers' search behaviour

	Av price	Price Dispersion
a	OK	✗
s	Nash predictions not rejected	Nash predictions not rejected

  - Av. Prices diminish with a , do not decrease with S
  - Price dispersion does not decrease with S.
- Similar results when increasing the memory of buyers from 50 to 150

Concluding remarks

- This talk: comparison between a Nash and an agent-based model: one standard case, two variations

	Individual Learning	Social Learning
FSS	(standard case) →	
Sequential Search	↓	(to be explored)

- Other dimensions (forth.) :
  - Fictious play or EWA
  - Turning to Individual data & dynamics

To open...

Two applied topics

Electronic or Retail Markets ?

Consumers learning on electronic markets

Shopbots

Thanks for your attention !

- **E-mail :** eric.darmon@univ-rennes1.fr
  roger.waldeck@enst-bretagne.fr

- **Executable & related material :** http://e.darmon.free.fr/fssmarket

References

- Related papers:
  - *"Can boundedly rational sellers learn to play Nash ? Pricing with adaptive sellers"*, E. Darmon and R. Waldeck, Journal of Economic Interaction and Coordination, 2006
  - "*Nash vs Reinforcement Learning on a Search Market: Some Similarities and Differences between Individual and Social learning*", E. Darmon and R. Waldeck, European Journal of Economic and Social Systems, *forth.*
  - *Search and Price Competition",* R. Waldeck, JEBO 2006

Nash

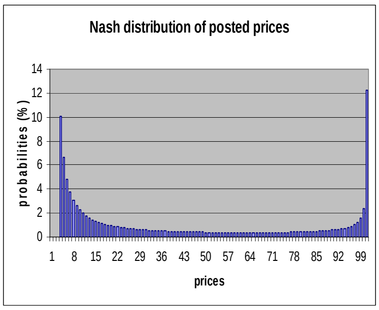

Learned

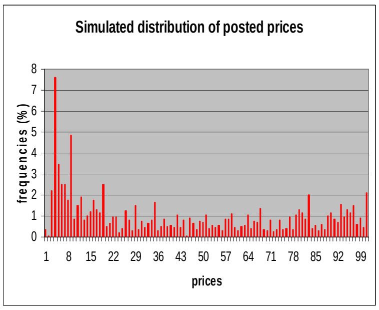

Parameter set : a = .9 ; k = 5 ; τ = 0.05 ; α = .8

Outline of the talk

- The model**s** : presentation & implementation
- The comparison in the reference case :
  - Main results
  - Behind the comparison: some inherent difficulties and limits
- Generalization & Validation issues
  - Robustness towards the learning algoritm
  - Robustness towards Buyers' search process
- Next …
- Related topics