Evolution and convergence of strategic behavior in continuous double auctions

Shira Fano, Marco LiCalzi, and Paolo Pellizzari

Ca' Foscari University of Venice, Italy

Evolution and market behavior in economics and finance
Pisa, 2–3 October 2009

Executive summary

In a call market, orders are cleared simultaneously.

Rustichini et al (1994) assume simple trading strategies.

They use analytical methods and characterize the equilibria.

The following results are well known.

As the market grows in size, allocative inefficiency tends to zero.

Equilibrium performance converges to the competitive outcome.

Evolution at work: equilibrium strategies morph from linear to flat functions of traders' strength.

Evolved behavior performs better than other "simple" equilibria.

Equilibrium strategies move towards truthtelling.
That is, agents become price-takers in large markets.

Executive summary

In a continuous double auction, orders are cleared asynchronously.

Rustichini et al (1994) assume simple trading strategies.

They use analytical methods and characterize the equilibria.

The following results are well known.

As the market grows in size, allocative inefficiency tends to zero.

Equilibrium performance converges to the competitive outcome.

Evolution at work: equilibrium strategies morph from linear to flat functions of traders' strength.

Evolved behavior performs better than other "simple" equilibria.

Equilibrium strategies move towards truthtelling.
That is, agents become price-takers in large markets.

Executive summary

In a continuous double auction, orders are cleared asynchronously.

We assume simple trading strategies.

They use analytical methods and characterize the equilibria.

The following results are well known.

As the market grows in size, allocative inefficiency tends to zero.

Equilibrium performance converges to the competitive outcome.

Evolution at work: equilibrium strategies morph from linear to flat functions of traders' strength.

Evolved behavior performs better than other "simple" equilibria.

Equilibrium strategies move towards truthtelling.
That is, agents become price-takers in large markets.

Executive summary

In a continuous double auction, orders are cleared asynchronously.

We assume simple trading strategies.

We use evolutionary methods to search for equilibria.

The following results are well known.

As the market grows in size, allocative inefficiency tends to zero.

Equilibrium performance converges to the competitive outcome.

Evolution at work: equilibrium strategies morph from linear to flat functions of traders' strength.

Evolved behavior performs better than other "simple" equilibria.

Equilibrium strategies move towards truthtelling.
That is, agents become price-takers in large markets.

Executive summary

In a continuous double auction, orders are cleared asynchronously.

We assume simple trading strategies.

We use evolutionary methods to search for equilibria.

We obtain the following results.

As the market grows in size, allocative inefficiency tends to zero.

Equilibrium performance converges to the competitive outcome.

Evolution at work: equilibrium strategies morph from linear to flat functions of traders' strength.

Evolved behavior performs better than other "simple" equilibria.

Equilibrium strategies move towards truthtelling.
That is, agents become price-takers in large markets.

Executive summary

In a continuous double auction, orders are cleared asynchronously.

We assume simple trading strategies.

We use evolutionary methods to search for equilibria.

We obtain the following results.

As the market grows in size, allocative inefficiency tends to zero.

Equilibrium performance converges to the competitive outcome.

Evolution at work: equilibrium strategies morph from linear to flat functions of traders' strength.

Evolved behavior performs better than other "simple" equilibria.

Equilibrium strategies move away from truthtelling.
That is, agents become price-fixers in large markets.

Executive summary

In a continuous double auction, orders are cleared asynchronously.

We assume simple trading strategies.

We use evolutionary methods to search for equilibria.

We obtain the following results.

As the market grows in size, allocative inefficiency tends to zero.

Equilibrium performance converges to the competitive outcome.

Evolution at work: equilibrium strategies morph from linear to flat functions of traders' strength.

Evolved behavior performs better than other "simple" equilibria.

Equilibrium strategies move away from truthtelling.
That is, agents become price-fixers in large markets.

Summary for executives

Leaves-eating animals abound.

When trees are low, gazelle do better.

When trees are high, giraffes do better.

Double auctions abound.

When orders are cleared simultaneously, price-taking does better.

When orders are cleared sequentially, price-fixing does better.

Summary for executives

Leaves-eating animals abound.

When trees are low, gazelle do better.

When trees are high, giraffes do better.

Double auctions abound.

When orders are cleared simultaneously, price-taking does better.

When orders are cleared sequentially, price-fixing does better.

Summary for executives

Leaves-eating animals abound.

When trees are low, gazelle do better.

When trees are high, giraffes do better.

Double auctions abound.

When orders are cleared simultaneously, price-taking does better.

When orders are cleared sequentially, price-fixing does better.

The environment

There is an an equal number n of buyers and sellers.

Each trader can exchange only one unit.

A buyer's valuation v and a seller's cost c are private information.
Assume they are i.i.d. on $[0, 1]$: the competitive price is $p^* = 1/2$.

If trade occurs at price p, payoffs are $v - p$ and $p - c$.

A *stronger* buyer has a higher valuation v.
A *stronger* seller has a lower cost c.

Stronger traders are more "valuable".
F.i., intramarginal traders are stronger than extramarginal traders.

Bilateral trading with simultaneous offers

There is only one buyer and one seller ($n = 1$).

Buyer and seller simultaneously submit a bid b and an ask a.
If $b \geq a$, trade occurs at the mid-price $p = (a + b)/2$.

The *bidding function* gives the buyer's bid $b = \beta(v)$.
The *asking function* gives the seller's ask $a = \alpha(c)$.

This model admits infinitely many equilibria.

One special case is the linear equilibrium

$$\beta(v) = \frac{2}{3}v + \frac{1}{12} \text{ if } v > \frac{1}{4} \quad \text{and} \quad \alpha(c) = \frac{2}{3}c + \frac{1}{4} \text{ if } c < \frac{3}{4}$$

Another special case is the constant equilibrium

$$\beta(v) = \frac{1}{2} \text{ if } v > \frac{1}{2} \quad \text{and} \quad \alpha(c) = \frac{1}{2} \text{ if } c < \frac{1}{2}$$

Allocative efficiency recommends the linear equilibrium.

Equilibrium strategies

Left: linear equilibrium. Right: constant equilibrium.

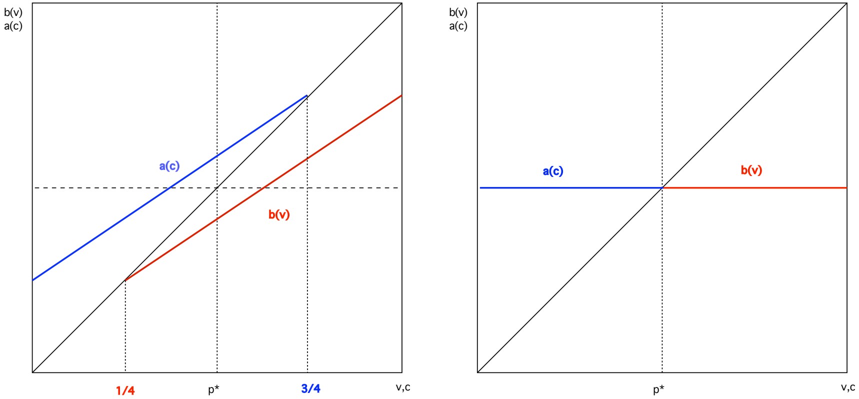

Individual rationality: $\beta(v) \leq v$ and $\alpha(c) \geq c$.

Truthtelling: $\beta(v) = v$ and $\alpha(c) = c$.

Weak traders' strategies do not matter, because they cannot trade.

Bilateral trading with sequential offers

There is one buyer and one seller ($n = 1$).

Buyer or seller arrive first, with equal probability.

Buyer and seller sequentially submit a bid b and an ask a.
If $b \geq a$, trade occurs at the earliest of the two prices.

The *bidding function* gives the buyer's bid $b = \beta(v)$.
The *asking function* gives the seller's ask $a = \alpha(c)$.

The equilibria are exactly the same as with simultaneous offers.

The only detectable difference is higher variability in the transaction price.

For $n = 1$, the two models start off with the same equilibria.

Bilateral trading with sequential offers

There is one buyer and one seller ($n = 1$).

Buyer or seller arrive first, with equal probability.

Buyer and seller sequentially submit a bid b and an ask a.
If $b \geq a$, trade occurs at the earliest of the two prices.

The *bidding function* gives the buyer's bid $b = \beta(v)$.
The *asking function* gives the seller's ask $a = \alpha(c)$.

The equilibria are exactly the same as with simultaneous offers.

The only detectable difference is higher variability in the transaction price.

For $n = 1$, the two models start off with the same equilibria.

The k-double auction

There are n buyers and n sellers.

Buyers and sellers submit price offers simultaneously.

Offers are aggregated into demand and supply functions.

Their intersection defines the market-clearing price p^*.

In any symmetric equilibrium, buyers shade their valuations and bid $\beta(v) < v$; similarly, sellers markup their costs and ask $\alpha(c) > c$.

This strategic misrepresentation departs from price-taking behavior.

When n gets large, the misrepresentation shrinks:
1) strategies converge to truthtelling: $\beta(v) = v$ and $\alpha(c) = c$;
2) the equilibrium price tends to the competitive price $p^* = 1/2$.

What if price offers arrive sequentially?

The k-double auction

There are n buyers and n sellers.

Buyers and sellers submit price offers simultaneously.

Offers are aggregated into demand and supply functions.

Their intersection defines the market-clearing price p^*.

In any symmetric equilibrium, buyers shade their valuations and bid $\beta(v) < v$; similarly, sellers markup their costs and ask $\alpha(c) > c$.

This strategic misrepresentation departs from price-taking behavior.

When n gets large, the misrepresentation shrinks:
1) strategies converge to truthtelling: $\beta(v) = v$ and $\alpha(c) = c$;
2) the equilibrium price tends to the competitive price $p^* = 1/2$.

What if price offers arrive sequentially?

The k-double auction

There are n buyers and n sellers.

Buyers and sellers submit price offers simultaneously.

Offers are aggregated into demand and supply functions.

Their intersection defines the market-clearing price p^*.

In any symmetric equilibrium, buyers shade their valuations and bid $\beta(v) < v$; similarly, sellers markup their costs and ask $\alpha(c) > c$.

This strategic misrepresentation departs from price-taking behavior.

When n gets large, the misrepresentation shrinks:
1) strategies converge to truthtelling: $\beta(v) = v$ and $\alpha(c) = c$;
2) the equilibrium price tends to the competitive price $p^* = 1/2$.

What if price offers arrive sequentially?

Our continuous double auction

There are n buyers and n sellers.

Agents arrive in random order and sequentially submit their offers.

Orders are immediately executed at the outstanding price if they are marketable; otherwise, they are recorded on the books.

When a transaction takes place, orders are removed from the book.

Orders are cleared asynchronously at different prices.

There are complications in moving from simultaneous to sequential offers. We simplify the strategy space.

The order of arrival of traders is randomly drawn over all queues.

The bid b_i of a buyer i is a limit order. (Same for the seller.)

When the buyer reaches the market, if the outstanding ask $a \leq b_i$, then his limit order is marketable and he trades at price $p = a$.

Otherwise, he records his limit order at b_i on the book.

Our continuous double auction

There are n buyers and n sellers.

Agents arrive in random order and sequentially submit their offers.

Orders are immediately executed at the outstanding price if they are marketable; otherwise, they are recorded on the books.

When a transaction takes place, orders are removed from the book.

Orders are cleared asynchronously at different prices.

There are complications in moving from simultaneous to sequential offers. We simplify the strategy space.

The order of arrival of traders is randomly drawn over all queues.

The bid b_i of a buyer i is a limit order. (Same for the seller.)

When the buyer reaches the market, if the outstanding ask $a \leq b_i$, then his limit order is marketable and he trades at price $p = a$.

Otherwise, he records his limit order at b_i on the book.

Simulation setup

We discretize types and strategies. Let δ be the tick size.

Buyers' valuation are drawn on the support $\{\delta, 2\delta, \ldots, (m-1)\delta\}$. (Same for sellers.)

The set of feasible offers is $O = \{\delta, 2\delta, \ldots, (m-1)\delta\}$.

We allow traders to play randomized strategies.

Traders interact over a number T of trading days.

We evolve the strategies using a genetic algorithm.

Learning is driven by the cumulated profit of the strategy, as encoded in the probability that a trader issues a given order.

The fitness of an agent is its cumulated profit over the past τ days.

Every τ days, agents update strategies using standard GA machinery (selection, crossover, and mutation).

Technically, we use a genetic algorithm known as "4-2".

Simulation setup

We discretize types and strategies. Let δ be the tick size.

Buyers' valuation are drawn on the support $\{\delta, 2\delta, \ldots, (m-1)\delta\}$. (Same for sellers.)

The set of feasible offers is $O = \{\delta, 2\delta, \ldots, (m-1)\delta\}$.

We allow traders to play randomized strategies.

Traders interact over a number T of trading days.

We evolve the strategies using a genetic algorithm.

Learning is driven by the cumulated profit of the strategy, as encoded in the probability that a trader issues a given order.

The fitness of an agent is its cumulated profit over the past τ days.

Every τ days, agents update strategies using standard GA machinery (selection, crossover, and mutation).

Technically, we use a genetic algorithm known as "4-2".

Transaction prices

R1. *The evolution of trading strategies stabilizes prices around the competitive price p^*.*

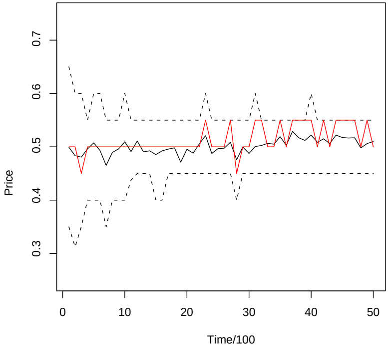

n	μ	σ
1	0.502	0.095
4	0.506	0.076
10	0.498	0.061
100	0.497	0.017

R2. *If $n_1 > n_2$, the distribution of the transaction price $P(n_2)$ is more diffuse than $P(n_1)$.*

Traders' profits

R3. *The evolution of trading strategies improves aggregate profits and is more beneficial for stronger traders.*

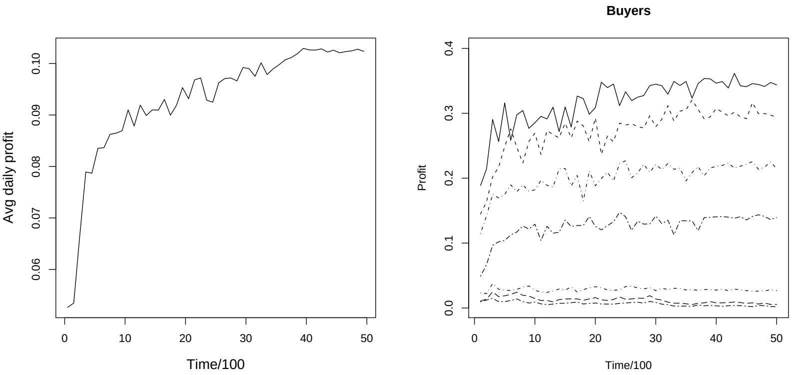

Figure: Left: traders' average daily profit. Right: time series of gains for buyers, arranged from top to bottom according to their strength.

Trading strategies

R4. *As n increases, the trading strategies of the intramarginal agents move away from linear functions of their strength towards a constant offer equal to the competitive price.*

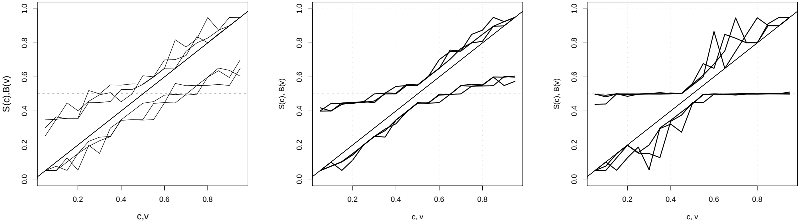

Extramarginal agents trade rarely and have nothing to learn.

The asymptotic behavior of the evolved strategies is exactly opposite to the k-double auction.

Incidentally, for $n = 1$ evolution selects the linear equilibrium.

Evolution "purifies" strategies.

We allow traders to evolve mixed strategies.

R5. *The evolution of intramarginal trading strategies leads to pure strategies.*

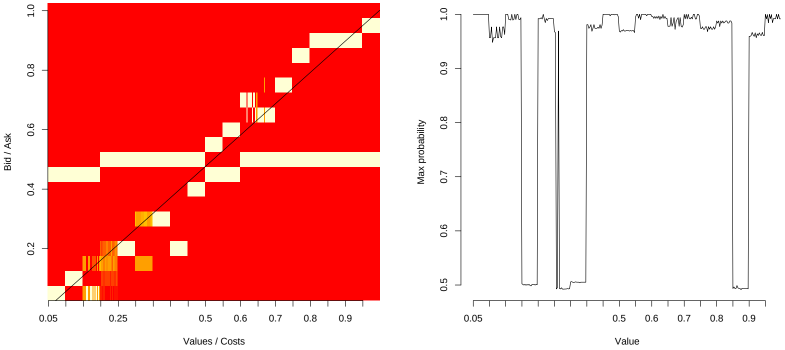

Left: mixed strategies for the whole population (red is zero). Right: modal probability distribution for each type of buyer when $n = 100$.

Robustness analysis

A natural benchmark is Zhan and Friedman (JEDC, 2007).

They search for *ex ante* equilibria within the class of *standard* markup strategies

$$[\text{ZF-}s] \qquad \beta(v) = v(1 - m_d) \quad \text{and} \quad a(c) = c(1 + m_u)$$

where $m_d, m_u \geq 0$ are the markdown and markup coefficients.

As shown elsewhere, symmetry requires a *convex markup*:

$$[\text{ZF-}c] \qquad \beta(v) = v(1 - m_d) \quad \text{and} \quad \alpha(c) = c + m_u(1 - c)$$

R6. *GA is evolutionarily stable against ZF-s or ZF-C. The converse is not true.*

Allocative efficiency

R7. *Allocative inefficiency is decreasing in market size.*

	continuous double auction			optimal	0.5-double auction	
n	GA	ZF-s	ZF-c	mechanism	least	most
1	0.199	0.202	0.371	0.16	0.16	1.00
10	0.084	0.115	0.131	0.004*	0.004*	0.004*
100	0.048	0.070	0.040	nil	nil	nil

Table: Allocative inefficiencies.

Distance from equilibrium

A strategy profile σ is an ε-*equilibrium* if, for all agents i and for all strategies s_i in S_i,

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i}) - \varepsilon \tag{1}$$

Clearly, σ is a Nash equilibrium iff (1) holds for $\epsilon = 0$.

Let $\varepsilon^*(\sigma) \geq 0$ the least ϵ that satisfies (1). We interpret is as the "distance" that separates σ from being an equilibrium.

Roughly speaking, $\varepsilon^*(\sigma)$ measures the worst case temptation of at least an agent to break away from σ.

	n	ε^*	type	offer	p-value
GA	1	0.00786	$v = 0.85$	0.65	0.03118
	10	0.00782	$v = 0.75$	0.55	0.00009
	100	0.00679	$v = 0.95$	0.50	0.01404
ZF-c	1	0.04170	$v = 0.90$	0.75	0.00000*
	10	0.02230	$v = 0.95$	0.60	0.00000*
	100	0.06492	$c = 0.10$	0.50	0.00000*

Evolutionary stability

GA invades ZF-c.

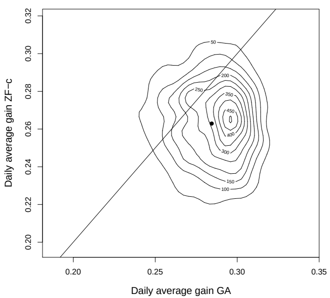

Figure: Left: joint distribution of daily gains for GA agents (x-axis) invading ZF-c population (y-axis). Right: incremental profits of GA buyers (versus ZF-c) as a function of their type.

ZF-c cannot invade GA.