Computation in a Distributed Information Market

Joan Feigenbaum	(Yale)
Lance Fortnow	(NEC Labs)
David Pennock	(Overture)
Rahul Sami	(Yale)

Markets Aggregate Information!

Evidence indicates that markets are good at combining information from many sources:

- Markets like the Iowa Electronic Market predict election outcomes better than opinion polls [Forsythe *et al.* ’99].
- Futures and options markets provide accurate forecasts of their underlying commodities/securities [Jackwerth *et al.* ’96].
- Sports betting markets provide unbiased forecasts of game outcomes [Gandar *et al.* ’98; Debnath *et al.* ’03]
- Laboratory experiments confirm information aggregation [Plott *et al* ’88, Plott *et al.* ’97]
- Markets sometimes deployed primarily for information aggregation (*e.g.*, IEM, Hollywood Stock Exchange)

Market as a Computation Device

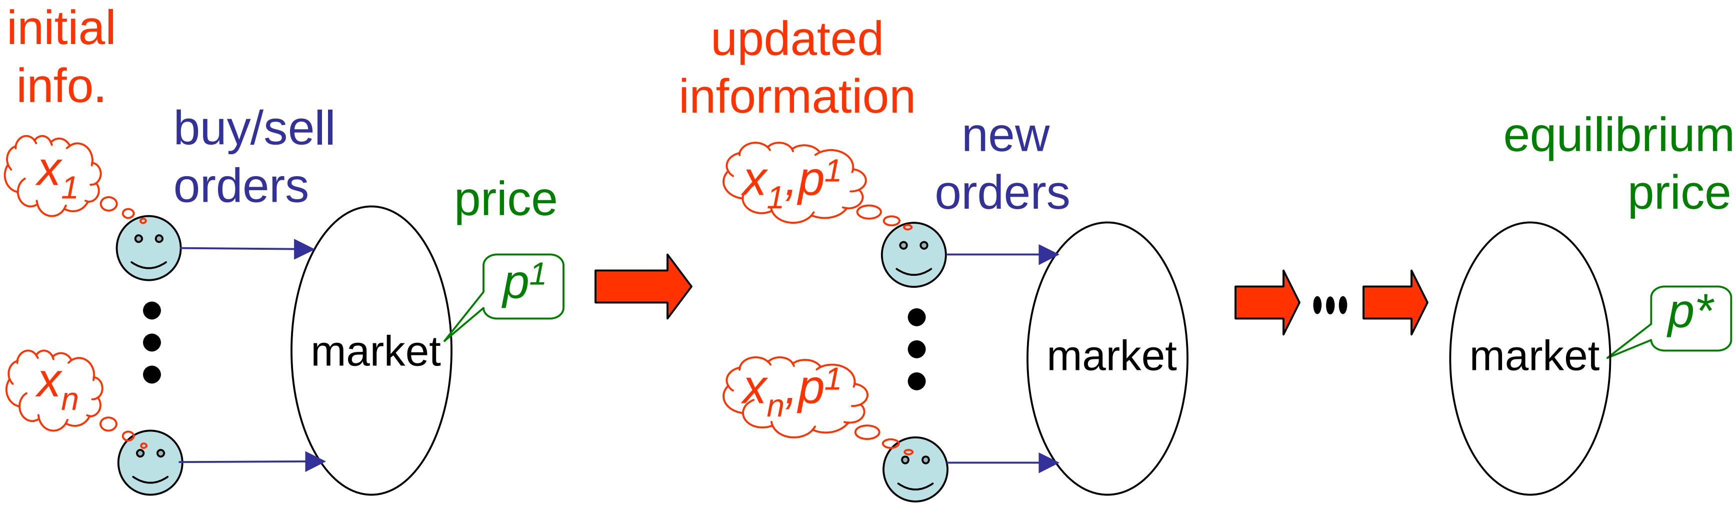

equilibrium price p^* $\equiv$ aggregate $f(x_1, x_2, \ldots, x_n)$

Questions:

- What aggregate functions *f ( )* can be computed?
- How many securities must be traded?
- How fast does the market price converge?

Simplified Market Model

- Study Boolean functions
  - Each trader i has a single bit of information x_i
  - Desired aggregate is a Boolean function $f(x_1, x_2, \ldots, x_n)$.
- Trade in a single security F with payoff contingent on f :
  If $f(x_1, x_2, \ldots, x_n)$ turns out to be 1, F eventually pays off \$1; otherwise, F eventually pays off \$0.
- Use *multiperiod Shapley-Shubik* model of the market
  - Trading occurs in a sequence of rounds.
  - In each round, trader i brings a "money supply" b_i and a "securities supply" q_i to the market.
  - Clearing price is $p = \Sigma b_i / \Sigma q_i$.
- Simplifying assumption: $q_i = 1$ (forced trade)
- Trader behavior: common-prior Bayesians, *truthful* (non-strategic) bidding.

Example: *OR* function

- Two traders, who initially know x_1, x_2 respectively. Uniform prior distribution on (0,0), (0,1), (1,0), (1,1).

- Single security *F*, based on $f(x_1,x_2) = x_1 \vee x_2$. *F* has value \$1 if $x_1 \vee x_2 = 1$; value \$0 otherwise.

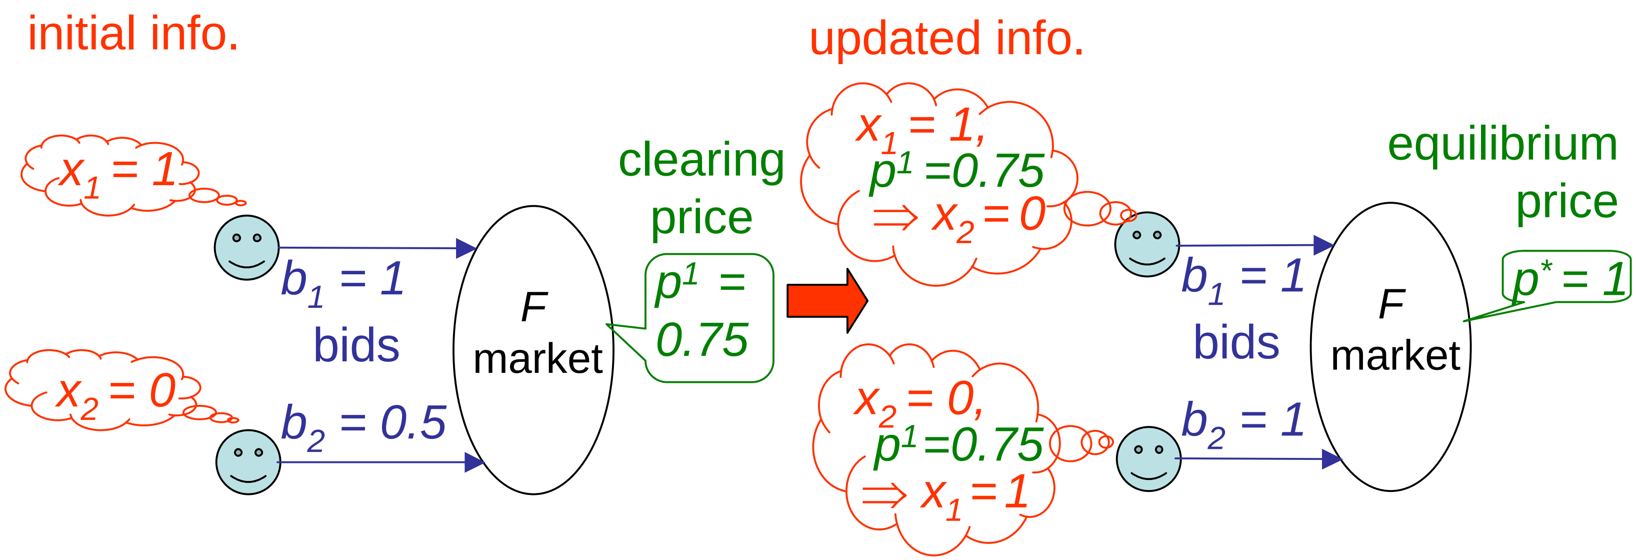

ü Equilibrium price reveals the value of $f(x_1,x_2)$ in this market.

Example: *XOR* function

- Two traders, who initially know x_1, x_2 respectively.
  Uniform prior distribution on (0,0), (0,1), (1,0), (1,1).
- Single security F, based on $f(x_1,x_2) = x_1 \oplus x_2$.
  F has value \$1 if $x_1 \oplus x_2 = 1$; value \$0 otherwise.

initial info.

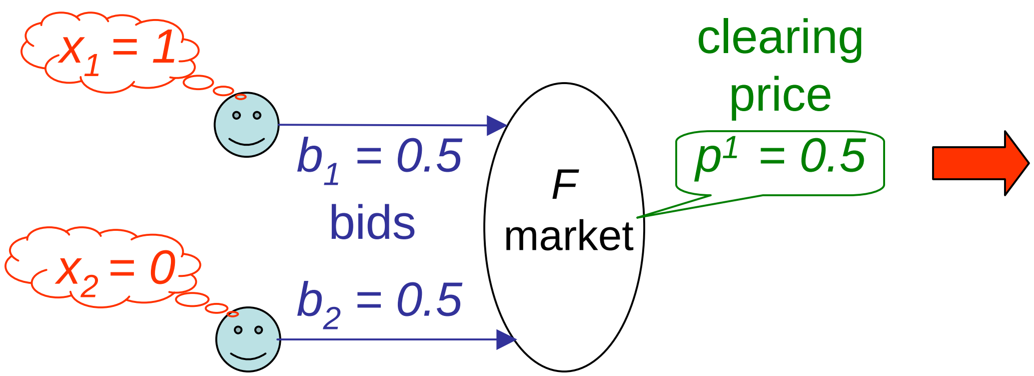

Clearing price reveals
no new information
to either agent
$\Rightarrow$ equilibrium reached!
Price $p^* = 0.5$.

X Equilibrium price does *not* reveal the value of $f(x_1,x_2)$ here.

Theorem: Computable Functions

If f can be expressed as a *weighted threshold function*

$$f(x_1, x_2, \ldots, x_n) = 1 \text{ if } \Sigma w_i x_i \geq 1$$
$$= 0 \text{ if } \Sigma w_i x_i < 1,$$

then, for *any* prior distribution, the market price of F converges to the true value of $f(x_1, x_2, \ldots, x_n)$.

e.g., *OR* function:

$$x_1 \vee x_2 \vee \ldots \vee x_n = 1 \text{ iff } \Sigma x_i \geq 1$$

Proof uses McKelvey-Page theorem on common knowledge of aggregates, combined with a counting argument.

Converse: Non-computable Functions

If f *cannot* be expressed as a weighted threshold function, then there exist prior distributions for which the price of F does not converge to the true value of $f(x_1, x_2, \ldots, x_n)$.

Proof idea:

- Construct a probability distribution over x such that

$$\forall i \;\; P(x_i = 1 \mid f(x) = 1) \;=\; P(x_i = 1 \mid f(x) = 0)$$

- By Bayes' law,

$$\forall i \;\; P(f(x) = 1 \mid x_i = 1) \;=\; P(f(x) = 1 \mid x_i = 0)$$

$\Rightarrow$ Each agent i's bid is independent of x_i!

Converse Proof (1)

- *f* is not a weighted threshold function

$\Rightarrow$ Convex hulls of $f^{-1}(0)$ and $f^{-1}(1)$ intersect in R^n.

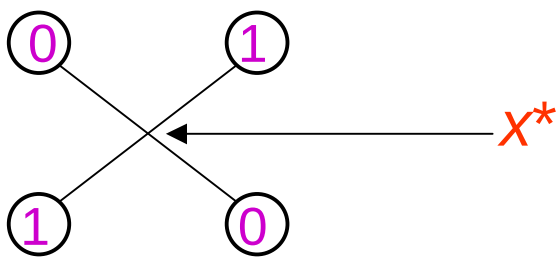

- There is a point x^* in intersection of $f^{-1}(0)$ and $f^{-1}(1)$.

Converse Proof (2)

- Can find probabilities $\lambda_1, \lambda_2, \ldots$ and $\mu_1, \mu_2, \ldots$ such that

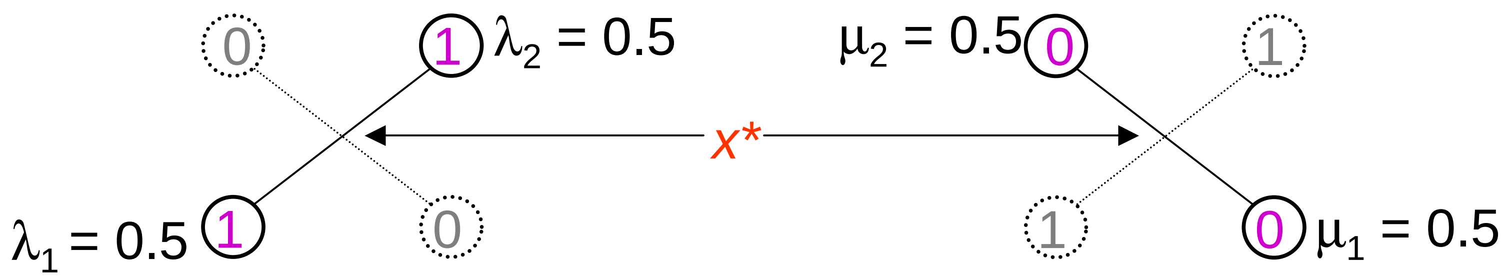

$E(x \mid f(x) = 1) = x^*$ $\qquad$ $E(x \mid f(x) = 0) = x^*$

- Take the mean of the two distributions:

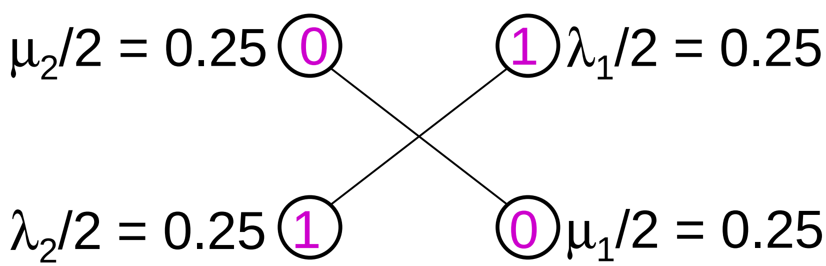

$$\forall i \;\; P(f(x)=1 \mid x_i=1) \; = \; P(f(x)=1 \mid x_i=0) \; = \; 0.5 \,!$$

Convergence Time Bounds

- Upper bound:
  For any function f, and any prior distribution, the market reaches the equilibrium price p^* after at most n rounds of trading.

- Lower bound:
  There is a family of weighted threshold functions C_n (the "carry-bit" functions) with $2n$ inputs, and corresponding prior distributions, such that it takes n rounds in the worst case to reach equilibrium.

Bounds are tight up to a factor of 2.

Lower Bound Proof (1)

- Function C_n: Carry bit of adding two n-bit numbers

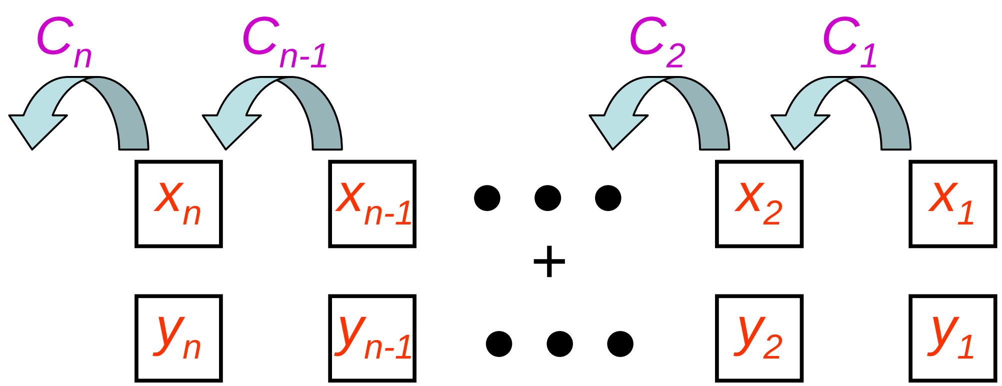

- Constructing the probability distribution P_n
  - (x_1, y_1) : uniformly distributed
  - (x_2, y_2) : distribution conditioned on (x_1, y_1), such that C_2 is independent of (x_1, y_1)
  - (x_i, y_i) : distribution conditioned on $(x_1, \ldots, x_{i-1}, y_1, \ldots, y_{i-1})$, such that C_n is independent of $(x_1, \ldots, x_{i-1}, y_1, \ldots, y_{i-1})$

Lower Bound Proof (2)

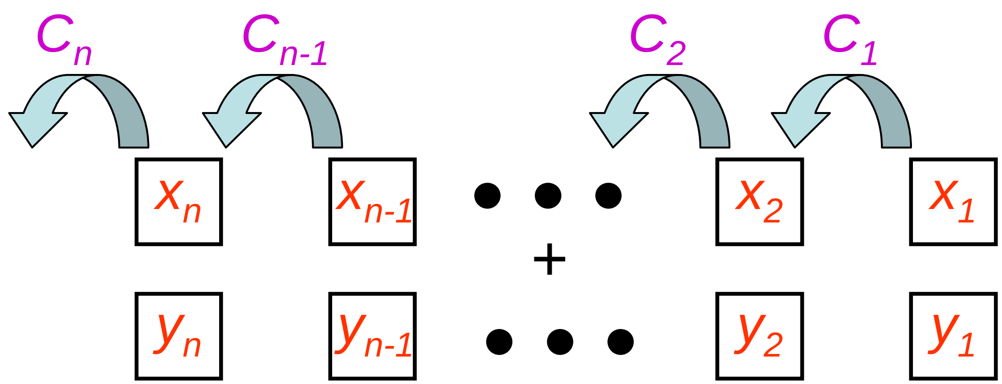

- First round: Except x_n and y_n, no other bit influences its owner's expectation of C_n.
  $\Rightarrow$ only $x_n + y_n$ revealed.

- If $x_n + y_n$ is revealed to be *1*, remaining problem is equivalent to computing C_{n-1}.

Summary

Analysis of a simple market model shows limits on what a market can compute and how fast it can compute.

Future directions:

→ More realistic market model
Strategic models, richer information, generic priors, complexity of traders' computation, *etc*.

→ Application to information market design
Design securities for certain convergence and faster convergence.