Computation in a Distributed Information Market

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

To appear in ACM EC ’03, San Diego, June 9-12, 2003.

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.* ’02]
- Laboratory experiments confirm information aggregation [Plott *et al.* ’88, Plott *et al.* ’97]
- Theoretical models: “Rational Expectation Hypothesis”

Market as a Computation Device

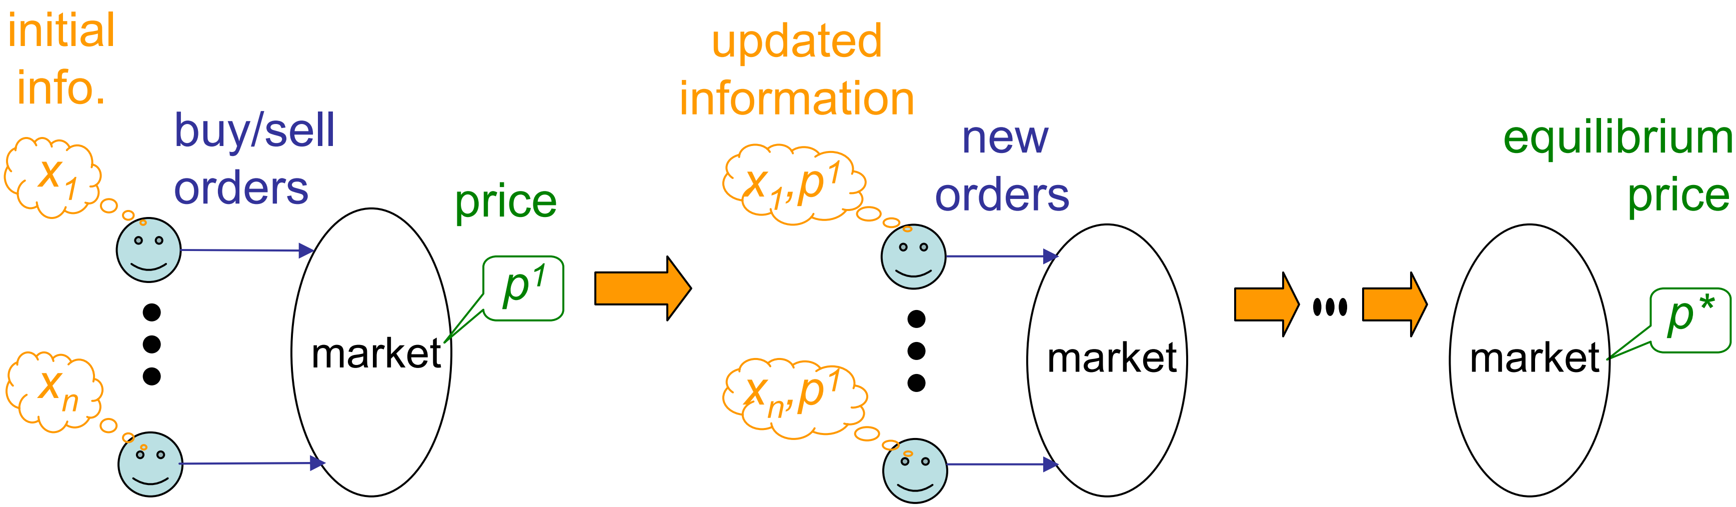

equilibrium price p^* $\equiv$ aggregate $f(x_1, x_2, \ldots, x_n)$

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: assume common-prior Bayesian, *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.

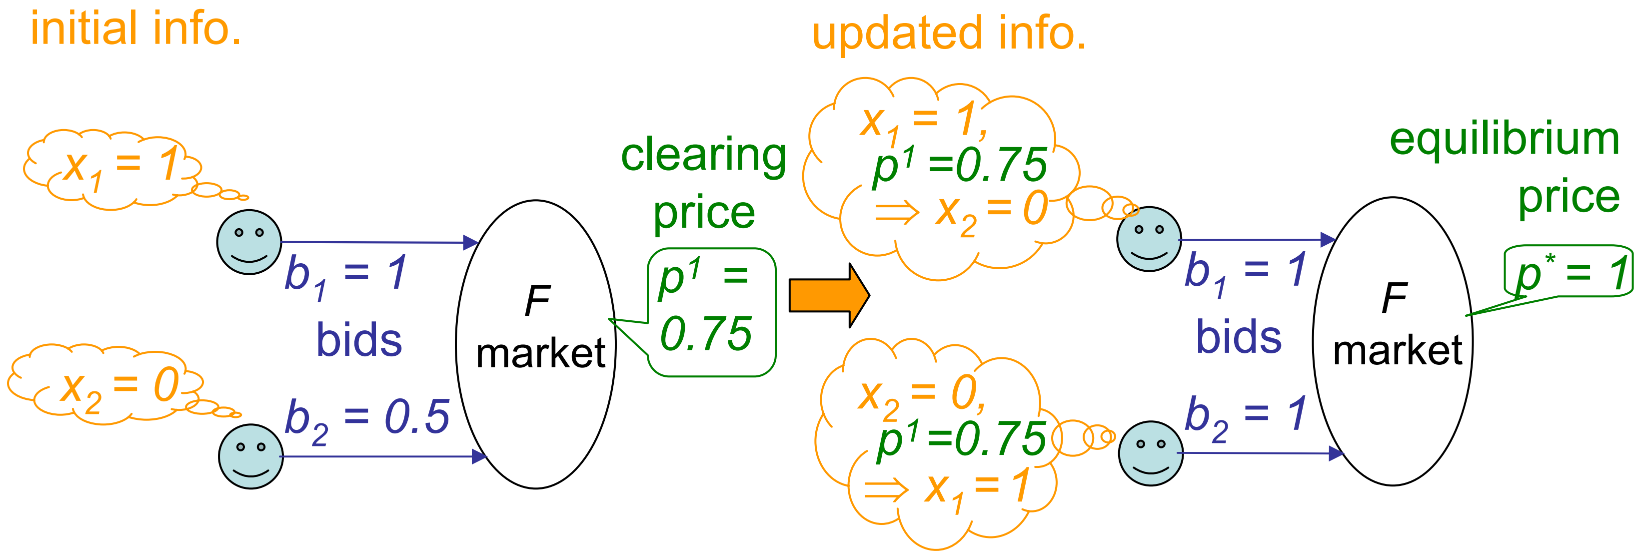

✓ Equilibrium price reveals the value of $f(x_1,x_2)$ in this market.

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.

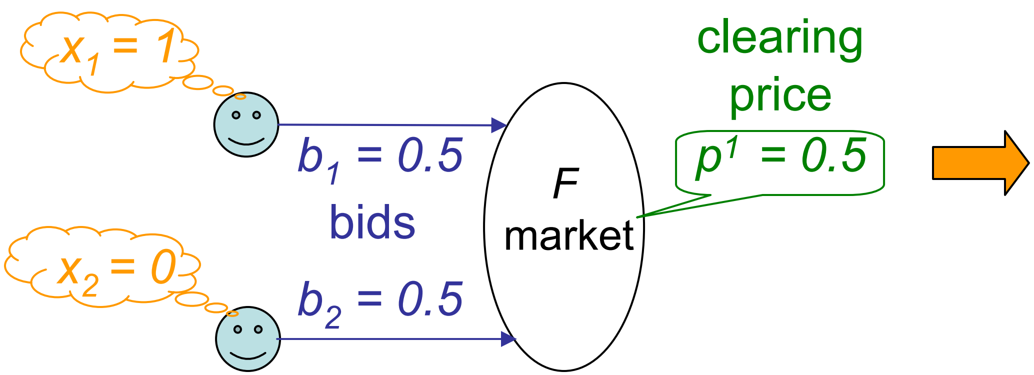

Clearing price reveals no new information to either agent
⇒ equilibrium reached!
Price $p^* = 0.5.$

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

Results: 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:

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

- 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 $f(x_1, x_2, \ldots, x_n)$.

e.g., X*OR* function.

Results: Convergence Time

- Upper bound:
  For any function *f* with *n* inputs, 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 2*n* 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.

Directions for Future Work

- Richer information: real numbers, distributions, *etc.*
- Effect of inaccurate or imprecise prices
- Strategic market models
- Can the common-prior assumption be relaxed?
- Complexity of *traders*' computations
- Distributed *computation*, e.g. through decentralized markets
- Convergence time for incremental updates
- Is there a connection to machine-learning theory?
- Designing information markets with multiple securities