Economic Efficiency Requires Interaction

Shahar Dobzinski, Sigal Oren, Noam Nisan

Overview

Sigal and Shahar's version

 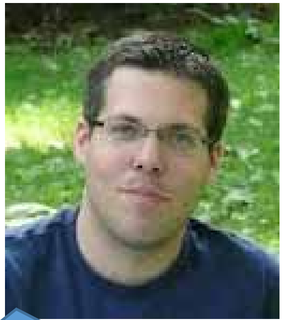

Here are some nice lower bounds on the simultaneous communication complexity of resource allocation.

Overview

Sigal and Shahar's version

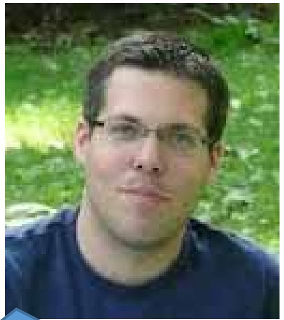

Here are some nice lower bounds on the simultaneous communication complexity of resource allocation.

My version

Here is a technical justification for Hayek's view that Markets are superior to Centralized Planning due to informational reasons.

Capitalism vs. Communism

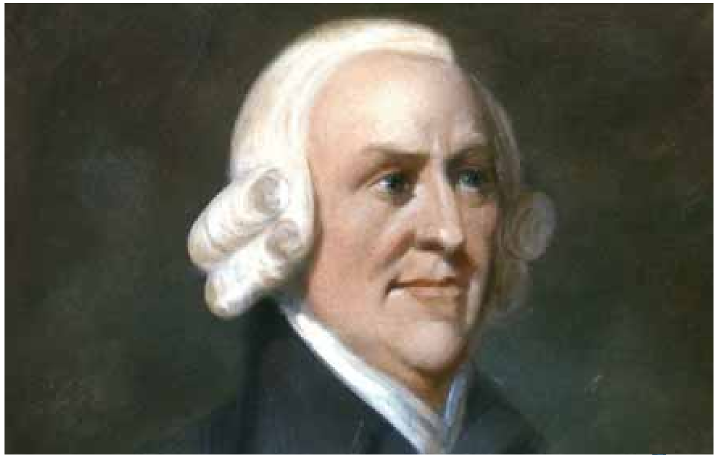

Adam Smith

It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest.

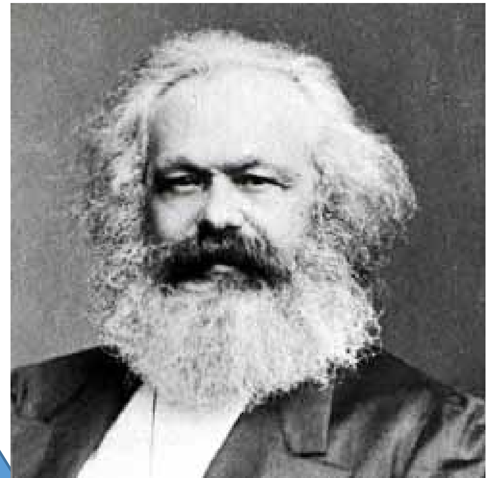

Karl Marx

The theory of Communism may be summed up in one sentence: Abolish all private property.

.

The Economic Calculation Debate

Socialism is unrealizable as an economic system because a socialist society would not have any possibility of resorting to economic calculation.

Ludwig von Mises

Input is Distributed

Which of these systems is likely to be more efficient depends (…) on whether we are more likely to succeed in putting at the disposal of a single central authority all the knowledge which ought to be used but which is initially dispersed among many different individuals, or in conveying to the individuals such additional knowledge as they need in order to enable them to fit their plans with those of others.

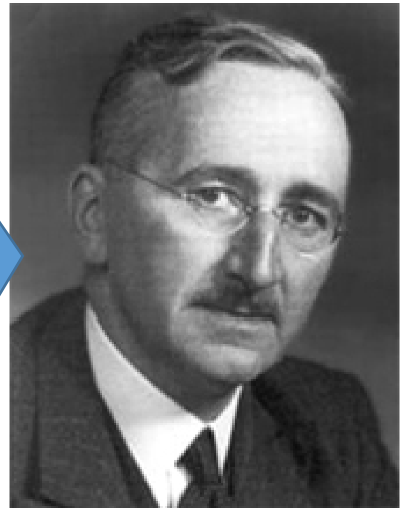

Friedrich A. Hayek

Formalizing Communication Protocols

The language of the command process is much larger than that of the Walrasian process. We must remember, however, that the pure command process is finished after only two exchanges of information while the tatonnement may go on for a long time

Leonid Hurwicz

Focus: Interaction

Markets

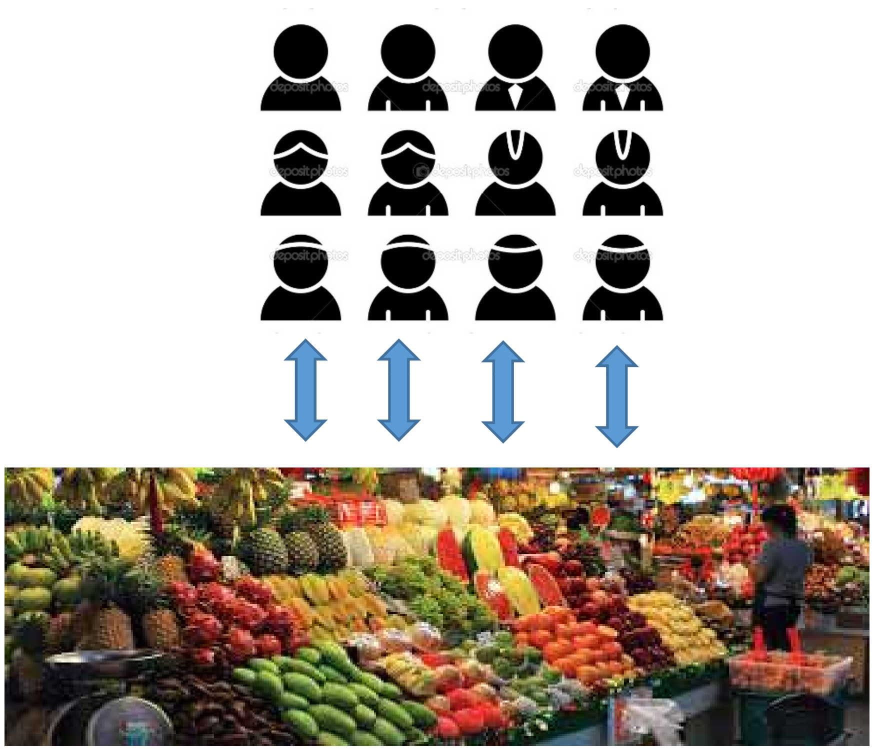

Centralized Planning

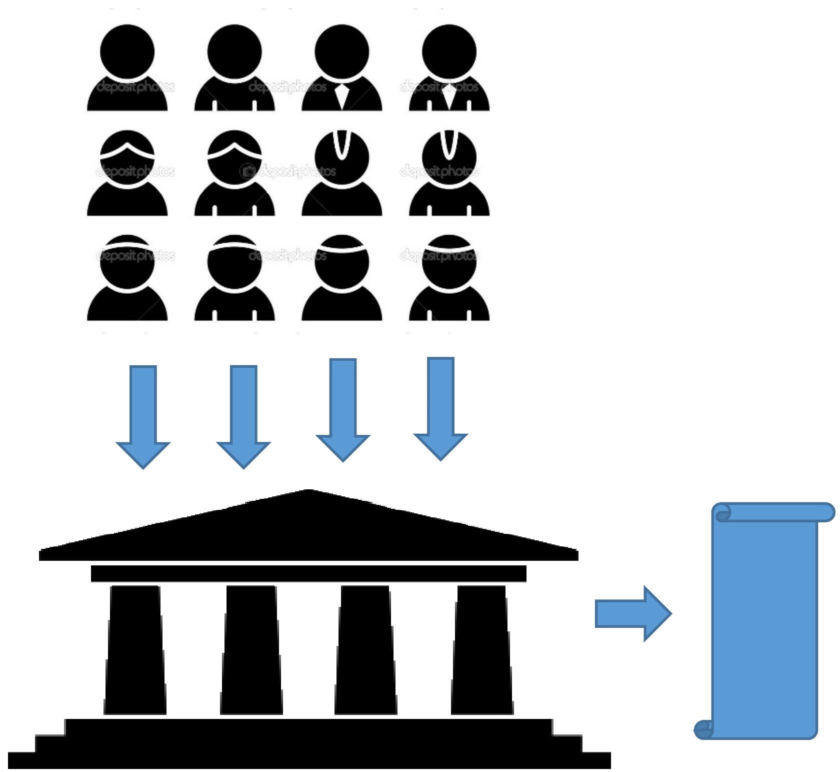

Our Results -- Overview

Combinatorial Auctions

THM: Interactive Protocols require exponentially less communication than do simultaneous ones to get a good approximation ratio with sub-additive valuations.

Unit Demand / Bipartite Matching

THM: Interactive protocols require exponentially less communication from each player to find an approximately maximum size matching.

Combinatorial Auctions

- A seller is selling *m* heterogeneous indivisible items to *n* buyers
- Each buyer *i* has a valuation function that assigns a real value to each subset of the items

- Valuations are:
  - Monotone: $v_i(S) \leq v_i(T)$ for $S \subseteq T$
  - Sub-additive: $v_i(S \cup T) \leq v_i(S) + v_i(T)$

- Goal: maximize social welfare $\Sigma_i\, v_i(S_i)$

Communication Model

- Initially each player *i* only knows his own sub-additive valuation $v_i()$
- They exchange messages according to a fixed protocol
  - Allow "broadcast", equivalently a "blackboard", equivalently a "coordinator"

Known: Finding the optimal allocation requires exponential communication. Nisan-Segal

Known: A constant factor approximation can be computed with polynomial communication. Feige

Can a *simultaneous* protocol approximate welfare well with polynomial communication?

Probably Yes

Selling the items simultaneously and separately gives a constant factor approximation at equilibrium (for any Bayesian prior on each player's valuation)

Christodoulou et al … Feldman et al

Missing: Correlated Bayesian prior

Probably No

Any polynomial sketch of sub-additive valuations has an approximation error of $\sqrt{m}$

Balcan et al ,Badanidiyuru et al

Missing: Even to approximate welfare in allocation

Our Results for Combinatorial Auctions

Theorem: Every simultaneous protocol requires exponential communication to get a better than $m^{1/4}$ approximation. (Even for XOS valuations)

Theorem: There exists a simultaneous protocol that uses polynomial communication to get a $m^{1/3}$ approximation.

Theorem: There exists a *polylog*-round protocol that uses polynomial communication to get a *polylog* approximation.

Bipartite Matching (unit demand bidders)

- *n* players, *n* items
- Each player *i* knows the subset S_i of items that he is interested in.
- Goal: find an (approximately) maximum size matching
- Communication goal: polylog bits from each player

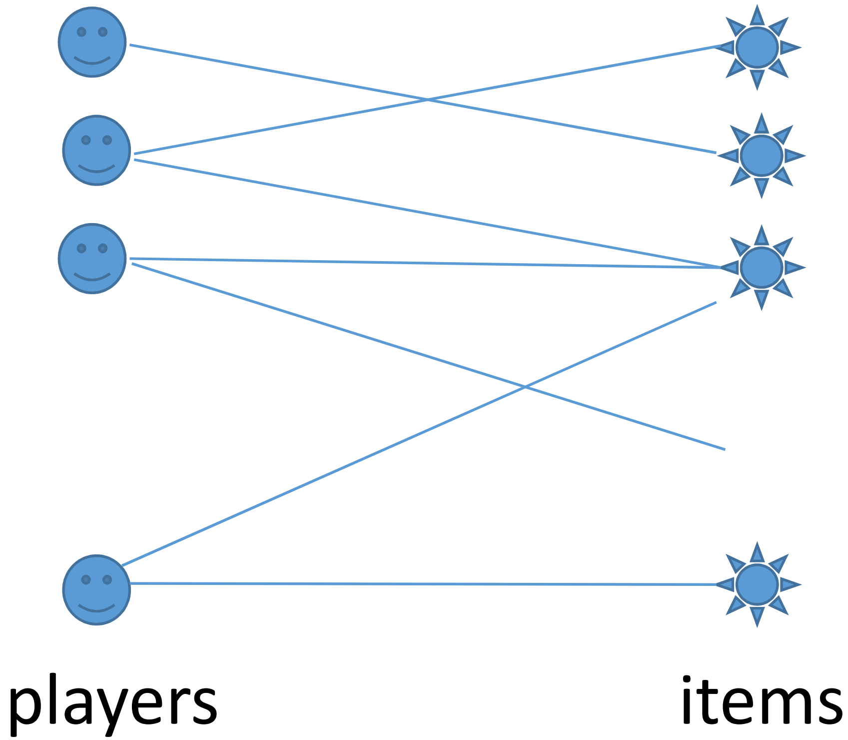

Auction-like Bipartite Matching

Demange-Gale-Sotomayor, Bertsekas

- Initialize all item prices to 0
- Repeat
  - Take any player *i* that is not holding an item
  - *i* takes the cheapest item that he wants (from whoever currently holds it)
  - The price of this item is increased by δ
- Until no unallocated player wants an item that is priced less than 1

Known: After n/δ steps (each requiring *log n* bits of communication) reaches a $(1-\delta)$-approximate maximum matching.

Question: How well can a simultaneous protocol do?

Our Results for Bipartite Matching

Theorem: Any deterministic simultaneous protocols using n^{ε} communication from each player can get only a $n^{1-\varepsilon}$ approximation.

Theorem: Any randomized simultaneous protocols using n^{ε} communication from each player can get only a $n^{1/2-2\varepsilon}$ approximation. (Implied also by Huang et al)

Theorem: A randomized simultaneous protocols using *O(log n)* communication from each player can get a $n^{1/2}$ approximation.

Theorem: A randomized $(\log n/\delta^2)$-round protocol using *O(log n) bits per round* from each player can get a $(1-\delta)$-approximation.

Fast Parallel Convergence

- Initialize all item prices to 0
- Repeat
  - **Every** player *i* that is not holding an item takes a **random** lowest-price item that he wants (simultaneously; collisions are dealt with arbitrarily)
  - The price of every item that changed hands is increased by δ
- Until no unallocated player wants an item that is priced less than 1

Theorem: After expected $\log n/\delta^2$ rounds (each using $\log n$ bits of communication from each player) reaches a $(1-\delta)$-approximate maximum matching.

Hard Distribution for Simultaneous Protocols

- *√n* popular items; each player wants a random half of them
- Each player also wants a single other random item (➔ expected linear size matching)
- Players don't know which of their items is non-popular
- |Matching| ≤

  √n + |Matched non-popular|

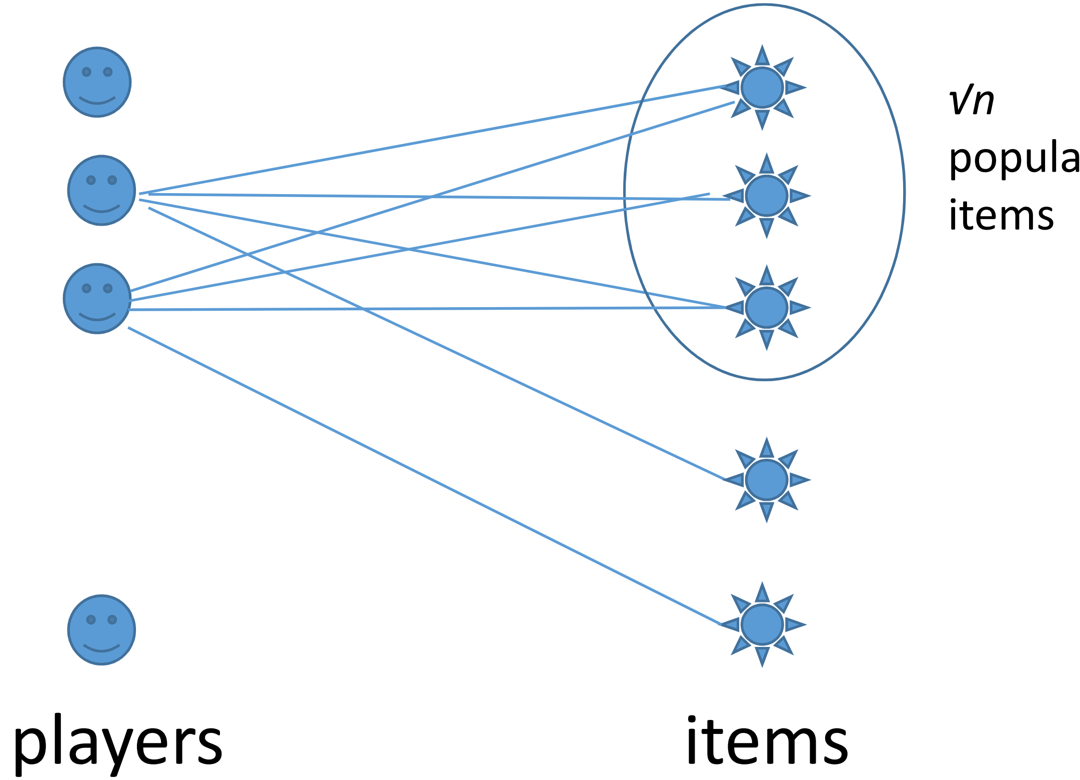

Reduction to 2-player Disjointess

- $|S_i|$ = items player i likes
- $|T_i|$= Non-popular items
- Chosen at random with $|S_i \cap T_i|=1$

- If you can find the item then you can distinguish from $|S_i \cap T_i|=0$ ➔ can't find it

Razborov

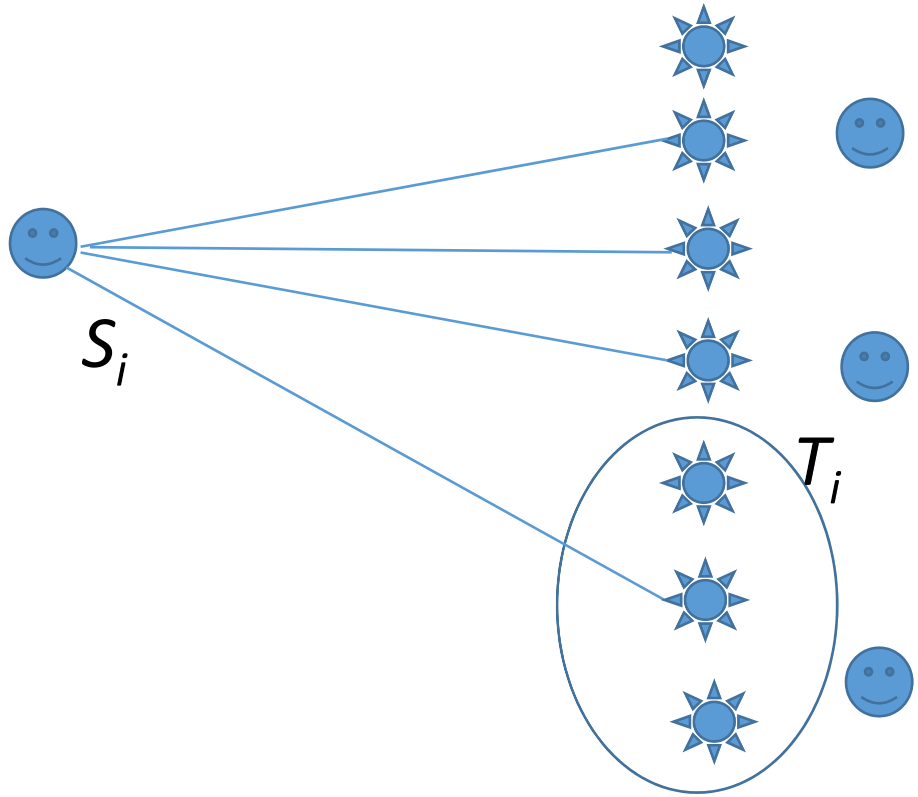

Open Problems

- Simultaneous protocols for combinatorial auctions with sub-modular (or Gross Substitute) bidders?
- Fast (polylog) approximate convergence of markets beyond bi-partite matching (general unit-demand or even Gross Substitutes)?
- Communication lower bounds for (even exact) bipartite matching?
  - Lower bound *nlogn;* upper bound $n^{1.5}$
  - Approach for understanding algorithmic complexity of matching

Thank You

Backup

Distributed vs. Centralized

Centralized can simulate Distributed

- (Revelation Principle: also incentives)

“Market Socialism” model:

- Central Planning Board + Trial-and-error price adjustments

Promises Professor von Mises a statue in the marble halls of the future Central Planning Board

Oskar R. Lange

Abba Le