Incentive-compatible Social Choice

Boi Faltings

Swiss Federal Institute of Technology (EPFL)

Social Choice

Choose an outcome o ∈ {o1,...} that a set of agents $A_1,..,A_k$ agree on

Examples:

- How to share airspace, radio spectrum, power lines, etc.
- Public policy decisions
- Dividing an inheritance
- ...

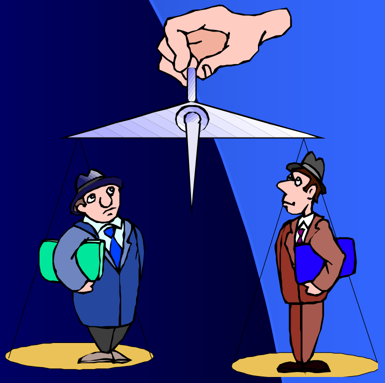

Social Choice Mechanisms

Mechanism:

protocol that chooses an outcome

Goals:

- Efficiency: maximize agents' combined utility (payoff)
- Individual rationality: each agent is better off participating than not

Commonly used Mechanisms

- Round-robin:
  ignores differences
- Lotteries:
  suboptimal outcome
- Voting:
  limited to simple decisions

Can we do better with computers?

Outline

- Example: slot allocation
- Optimizing social choice
- A truly incentive-compatible mechanism
- Incentive-compatible coordination

Example: Slot Allocation

- Airport runway can take only 1 plane/minute: 1 slot/minute
- Allocation:
  - what slot requests are granted?
- Coordination constraints:
  - flight needs takeoff and landing slot at different airports
  - flights need to be in sequence (rotation)

Current System

Slots are free

Allocation (each airport):

- 75% : historical rights to reuse same slots
- 25%: given on request in round-robin fashion, taking into account airline size

Coordination:

- IATA scheduling conferences (2/year)

Does it work?

- Demand exceeds capacity by up to 40%, but...
- 25% of allocated slots unused!

- Lack of coordination:

=> Airlines do no fly the best schedules

=> Poor service to the public!

Outline

- Example: slot allocation
- *Improving social choice*
- A truly incentive-compatible mechanism
- Incentive-compatible coordination

Improving Social Choice

- Optimization: choose allocation to optimize economic value
- Coordination: allocate matching pairs/sequences of slots

=> multi-agent problem-solving

Agent-based Social Choice (Coordination)

Preference/Constraint Elicitation

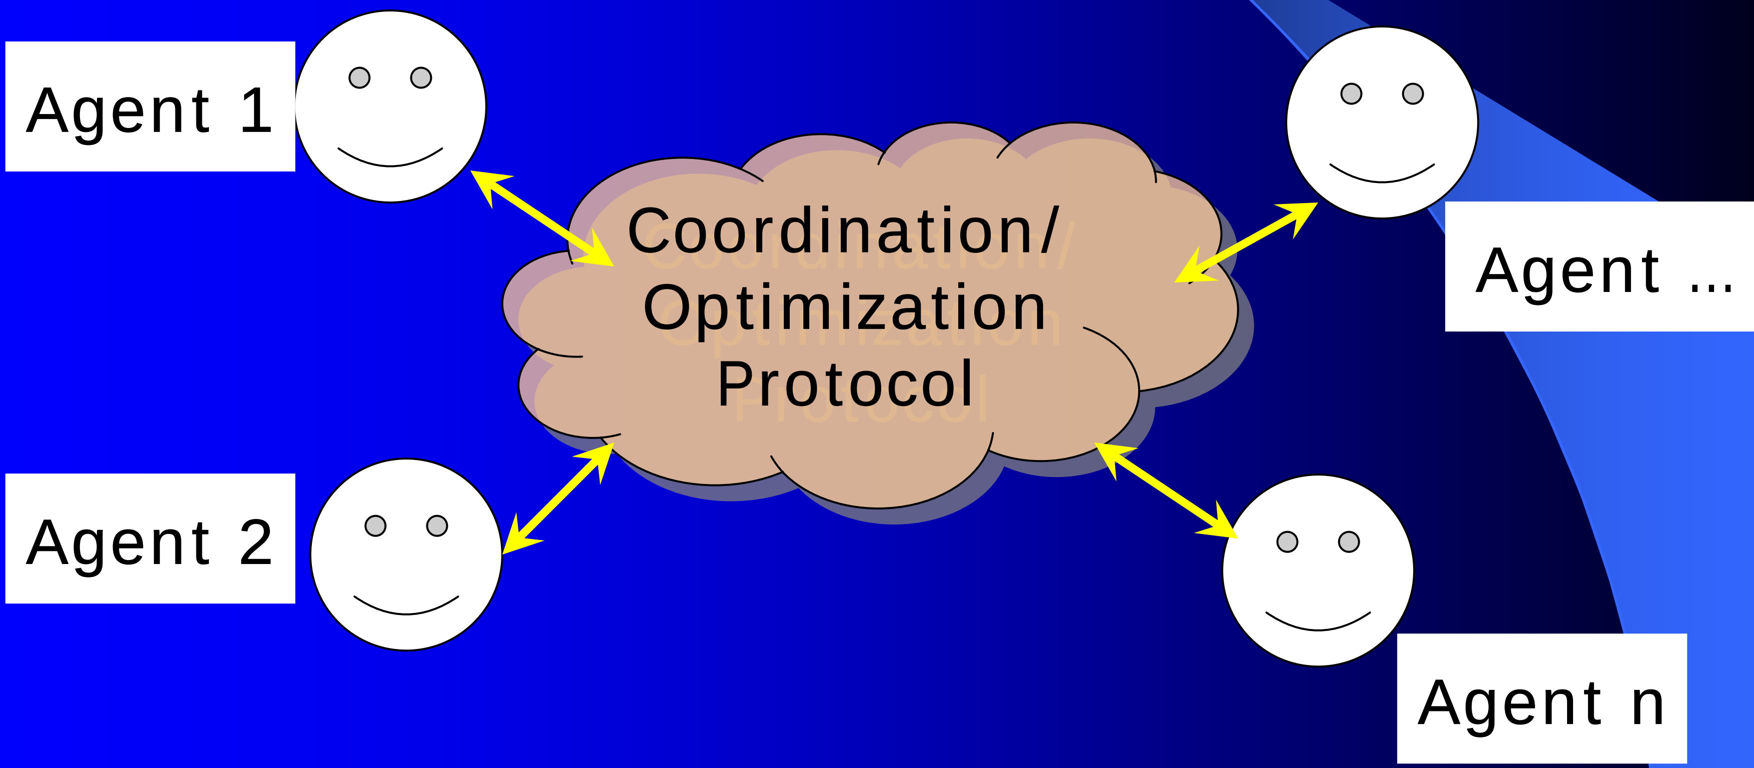

EPFL results...

Distributed Constraint Satisfaction:

- AAS + successors (Silaghi, 2000)
- Breakout (Eisenberg, Petcu, 2003)

Preference/Constraint Elicitation:

- Open Constraint Programming (Macho-Gonzalez, 2002)
- Example-critiquing (Torrens, Pu, ..., 1997-2004)

Applications:

- Communication networks (Willmott, Calisti,1998-2001)
- Project coordination (Eisenberg,2002)

Incentive-compatibility

Agents have conflicting incentives
=> do not cooperate for best solution

Social choice mechanism should make incentives *compatible*:
=> agents are best off by cooperating

Simple social choice

- Agents communicate their utility for each outcome
- Algorithm chooses outcome that maximizes sum of the utilities
- Centralized or distributed protocols

Example

- Airport has 2 slots
- 4 airlines A_1-A_4 want to use a slot
- They value its utility as follows:

A_1	A_2	A_3	A_4
10	8	3	1

Choosing a solution

Maximize sum of values:

- A_1 and A_2 get the slots

=> A_3, A_4 would exaggerate their utilities!

=> coordination and optimization make no sense!

Auctions

- Charge a variable fee for each slot
- English auction: increase fee until demand = supply

A_1	A_2	A_3	A_4
10	8	3	1

$\Rightarrow A_1$, A_2 can fly; each pays \$3(+ε)

Give revenue (\$6) to airport

Incentive-compatibility

IC mechanism makes equivalent:

- optimizing agent’s own utility
- optimizing combined utility

Auction achieves IC for airlines:

only agents with highest valuations have interest in winning auction

Incentives

If a runway is “closed for maintenance”:

- only A_1 gets a slot
- how much does it pay?

A_1	A_2	A_3	A_4
10	8	3	1

\$8(+ε) > \$6: airport revenue increases!

=> bad service is rewarded

Incentive-compatibility

- Auction is incentive-compatible for airlines, but not for airports!

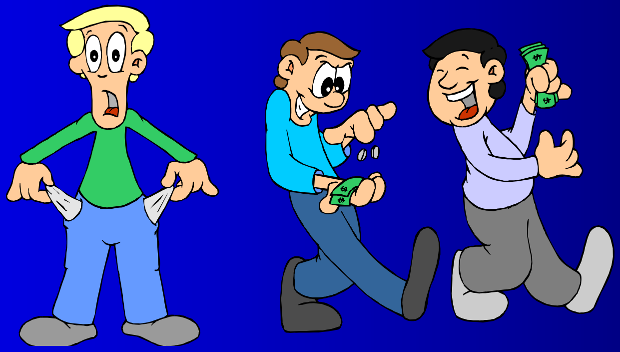

- Right incentive: airport has no revenue from auction, but only from fees

An Impossible Objective?

Game theory:

impossible to simultaneously have:

- Budget-balance (no revenue/loss)
- Incentive compatibility
- Individual rationality
- Efficiency (optimality)

Proposals

- DaGVA: return to each agent its *expected* payment (known in advance) [d'Aspremont & Gerard-Varet, 1979]
- Automated mechanism design: design a mechanism for a specific situation [Sandholm, 2003]
- Approximate IC, etc... [Parkes et al., 2001]

Outline

- Example: slot allocation
- Optimizing social choice
- *A truly incentive-compatible mechanism*
- Incentive-compatible coordination

Revenue-free Auctions

Solution: give up optimality

- choose one agent to be excluded
- auction slots among remaining agents
- give revenue to excluded agent
- excluded agent chosen independently of declarations (random, round-robin, etc.)

Example

A_1 excluded => valuations:

A_1	A_2	A_3	A_4
10	8	3	1

$\Rightarrow A_2$, A_3 get a slot; each pays \$1(+ε)

Give revenue (\$2) to A_1

Example (2)

Left out	Winners	Payment
A_1	A_2,A_3	2*$1
A_2	A_1,A_3	2*$1
A_3	A_1,A_2	2*$1
A_4	A_1,A_2	2*$3

Expected Outcomes

Airline	P(slot)	E[Payment]
A_1	¾	$ ¾
A_2	¾	$ ¾
A_3	½	0
A_4	0	- 2*$ ¾

Assumption: each agent left out with p=1/4

Properties

- Incentive-compatible for airlines:
  - A excluded: declarations do not matter
  - A included: equal to auction
- Individually rational for airlines:
  - A excluded: receives payment
  - A included: equal to auction
- Incentive-compatible for airport:
  - Best service optimizes income

Properties (2)

- Solution is suboptimal:
  - E[Utility] = 15 instead of 18
- But auctions not optimal either:
  - Total airline utility = 18-6 = 12
- Utility almost always better than auctions!

Outline

- Example: slot allocation
- Optimizing social choice
- A truly incentive-compatible mechanism
- *Incentive-compatible coordination*

Coordination

- Allocation of available slots: variables $x_1,..,x_n$
- Only certain combinations are usable: constraints $c_i(x_j,x_k,..)$
- Slot combinations have commercial value: relations $r_i(x_j,x_k,..)$

Formalizing Social Choice

Constraint optimization problem (COP) <X,D,C,R>:

- X = set of n variables (choices)
- D = set of n domains (options)
- C = set of m constraints (restrictions)
- R = set of p relations (valuations)
- Relations belong to agents $A_1,..,A_k$:
  $R_i = R(A_i)$, $R = \cup R_i$

Efficient Solution

$V^{*}_{R}(X)$ = assignment to X that

- satisfies all constraints
- maximizes sum of utilities in R

Incentive-compatibility...

- "Auction" mechanism => VCG tax:

  $Pay(A_i) = \sum_{j \neq i} R_j(V^*_{R \setminus R_i}) - R_j(V^*_R)$

  ("damage" to others)
- Incentive-compatible: agents best off to declare their true relations

Revenue-free VCG Tax

As in revenue-free auction:

- Choose excluded agent A_e
- Others optimize outcome
- Pay VCG tax to excluded agent:

  $Pay(A_i\text{->}A_e) =$

$$\sum_{j \neq i,e} R_j(V^*_{R \backslash (R_i \cup R_e)}) - R_j(V^*_{R \backslash R_e})$$

Suboptimal solution, but how bad?

Evaluating Mechanisms

Methodology: evaluate average performance on randomly generated problem instances

- Structured: model a real-world scenario
- Unstructured: completely random

Resource allocation in networks

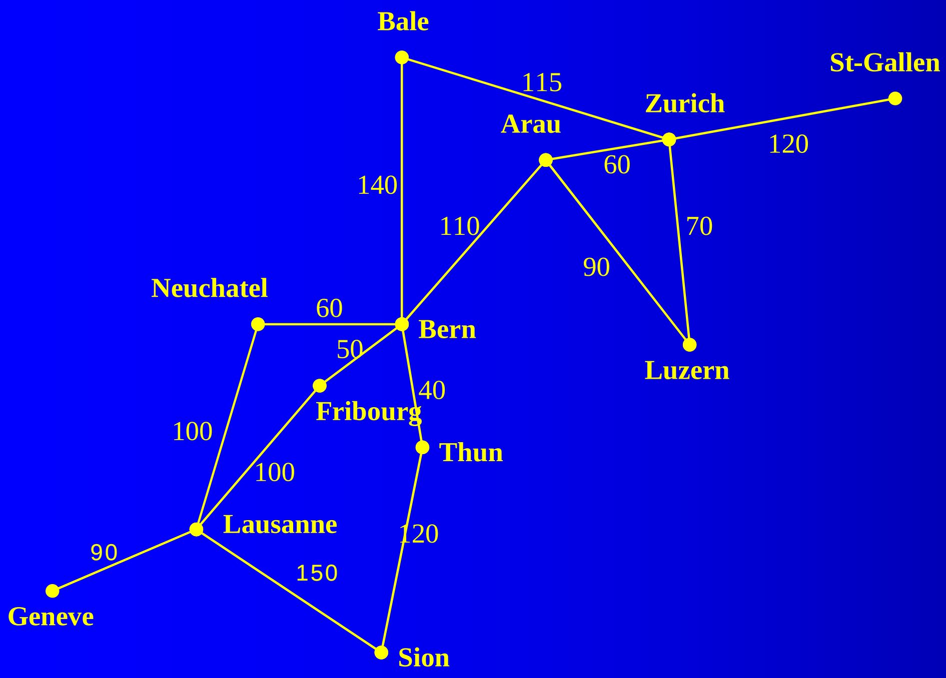

- Agents have different tasks and utilities
- Task = connect 2 nodes in graph
- Each link can only be used for one task

=> Allocate tasks to maximize revenue

Revenue-free vs. VCG

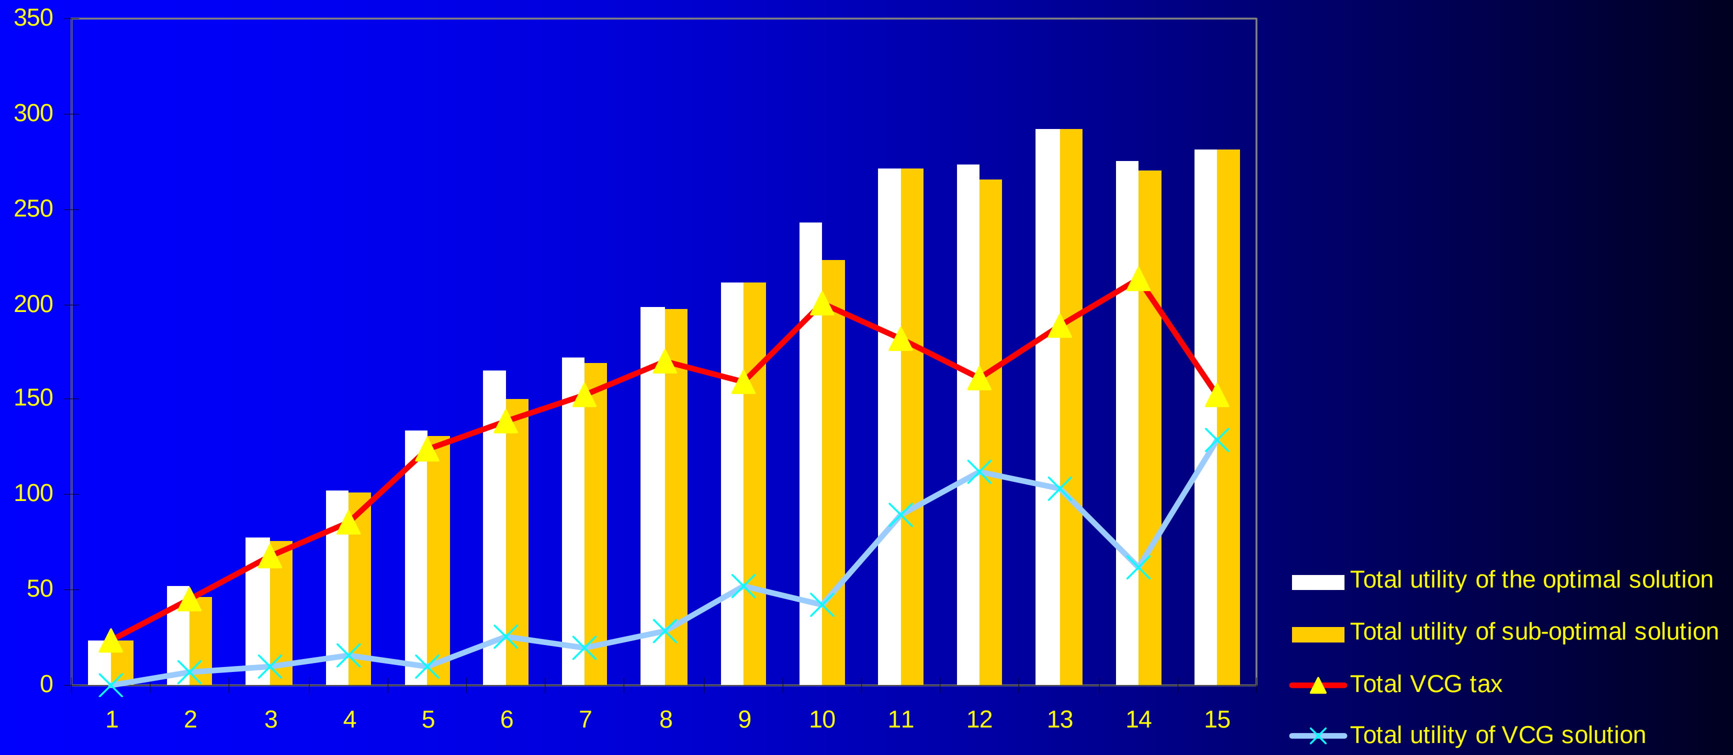

Unstructured problems

- Randomly generate set of variables, choices and constraints
- Relations = random value for each combination, uniformly distributed in [0..1], model cost
- Each agent seeks to minimize sum of its relations

Random Problems

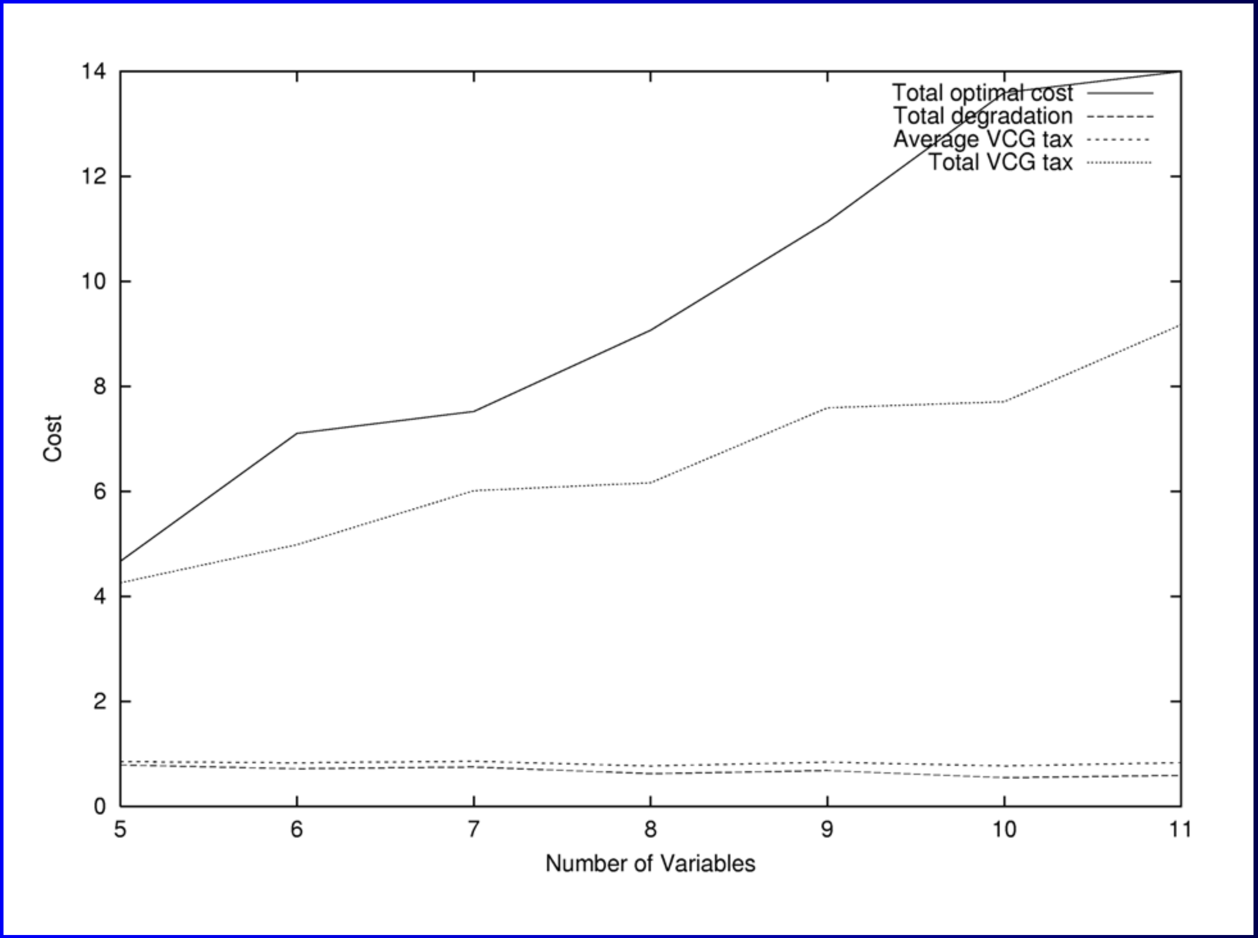

A problem...

One agent excluded everywhere

=> one airline gets no slots...

Solution:

- break relations into o subsets $R^1,..,R^o$ (orthogonal to agents)
- choose excluded agent $A_{e(j)}$ separately for each subset R^j
- optimize for $\bigcup_{j=1..o} R^j \setminus R^j_{e(j)}$

Distributing Exclusions

Tax = *average* tax if one agent (A_e) were excluded everywhere:

Pay(A_i->A_e) =

$$1/k \sum_{j\neq i,e} R_j(V^*_{R\setminus(R_i\cup R_e)}) - R_j(V^*_{R\setminus R_e})$$

Pay(A_e->A_i) =

$$1/k \sum_{j\neq i,e} R_j(V^*_{R\setminus(R_i\cup R_e)}) - R_j(V^*_{R\setminus R_i})$$

Pay(A_i) =

$$1/k \sum_{e\neq i} \sum_{j\neq i,e} R_j(V^*_{R\setminus R_i}) - R_j(V^*_{R\setminus R_e})$$

Distributing Exclusions

Relations have same distribution:

- expected influence identical for all agents
- to A_i, irrelevant whether A_j or A_l is excluded

=> For risk-neutral agents, mechanism remains individually rational and incentive-compatible

Generalizations

"Social compromise" (analogous to reverse auction):

- can be made IR/IC ex-post for risk-neutral agents

Double auctions:

- Can be made IR/IC ex-interim (in expectation) for risk-neutral agents

Implementation

- Required: optimization algorithm
- Centralized: integer programming, constraint programming
- Distributed: distributed constraint optimization
- Utility elicitation

=> open constraint programming

[Faltings and Macho-Gonzalez]

The future is bright...

- Agents allow complex protocols for efficient social choice
- Revenue-free mechanisms can align incentives for *all* participants
- => possible to greatly improve efficiency in many social choice problems

The Real “New Economy”

- Increasing population means increasing contention of resources

=> increasing need for social choice

- Traditional protocols are inefficient
- Market mechanisms create wrong incentives
- Agent-based systems can implement new decision mechanisms that provide the right incentives to everyone