Price of Anarchy in Auctions

Vasilis Syrgkanis
Microsoft Research NYC

Decentralization of Computer Systems and Services

- Large Scale Decentralized-Distributed Systems
- Multitude of Diverse Users with Different Objectives

Decentralization of Computer Systems and Services

- Large Scale Decentralized-Distributed Systems
- Multitude of Diverse Users with Different Objectives

Sharing Computing Resources

Decentralization of Computer Systems and Services

- Large Scale Decentralized-Distributed Systems
- Multitude of Diverse Users with Different Objectives

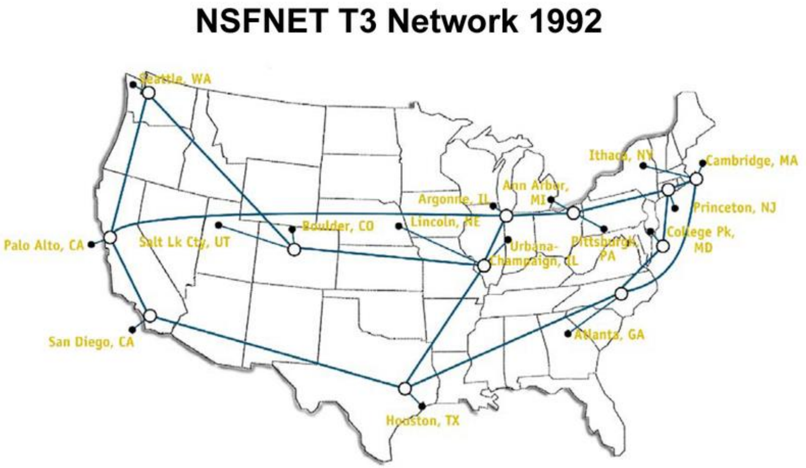

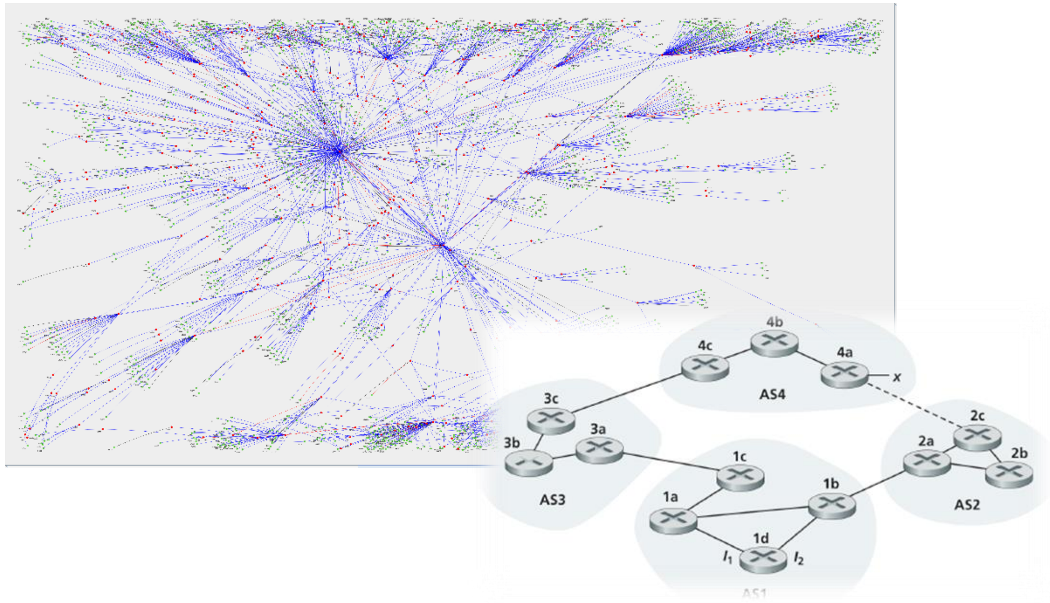

Network Routing

Decentralization of Computer Systems and Services

- Large Scale Decentralized-Distributed Systems
- Multitude of Diverse Users with Different Objectives

Online Advertising

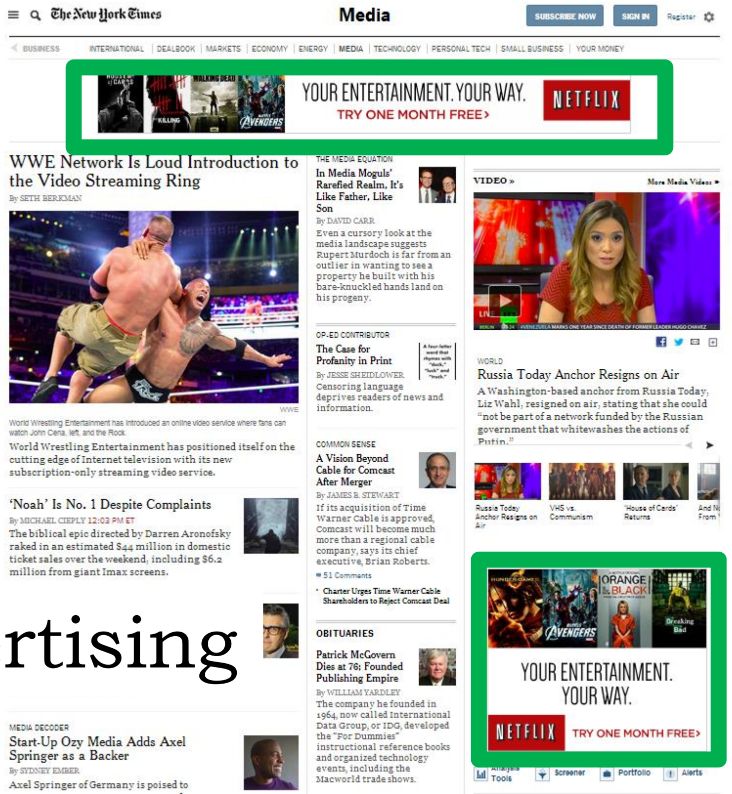

bing ads

ADECN

rightmedia

Google AdWords

double click

Common trend

Centralized, engineered systems with clear objectives → Platforms for interaction of diverse users

Strategic User Behavior

Centralized, engineered systems with clear objectives → Platforms for interaction of diverse users

Each optimizes their own objective

Strategic user behavior can cause inefficiencies.

Incentives and efficiency

Centralized, engineered systems with clear objectives → Platforms for interaction of diverse users

Analyze efficiency of systems taking into account strategic behavior of participants

Design systems for strategic users

Today: Electronic Markets

Ad Impressions

Goods

Information

Crowdsourcing

Today: Electronic Markets

Ad Impressions

Goods

Information

Crowdsourcing

Distinct properties of Electronic Markets

- Thousands of mechanisms run at the same time
- Players participate in many of them simultaneously or sequentially
- Environment too complex for optimal decision making
- Repeated game and learning behavior
- Incomplete information about environment (e.g. opponents)

Ad I

How should we design efficient mechanisms for such markets?

How efficient are existing mechanisms?

Information

Crowdsourcing

How should we design mechanisms such that a market composed of such mechanisms is approximately efficient?

We define the notion of a Smooth Mechanism.

A market composed of Smooth Mechanisms is globally approximately efficient at "equilibrium" even under learning behavior and incomplete information.

15

Example of strategic inefficiency

Efficient single-item auction

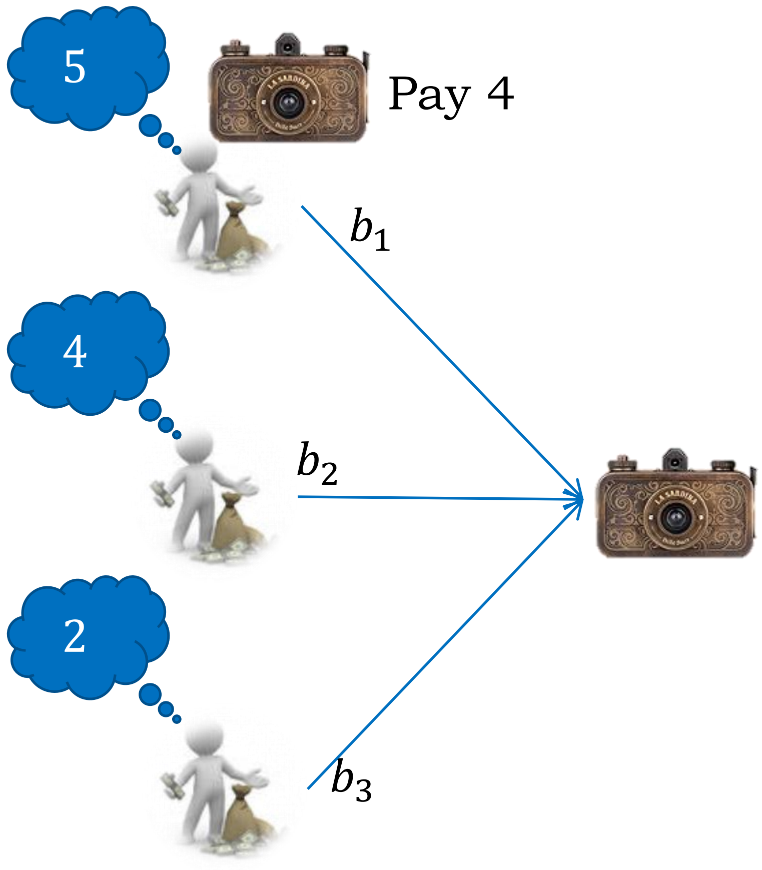

Vickrey (second-price) Auction

- Solicit bids
- Award to highest bidder
- Charge second highest bid

Classic Result. Dominant strategy equilibrium is efficient. Highest value wins

Two second price auctions

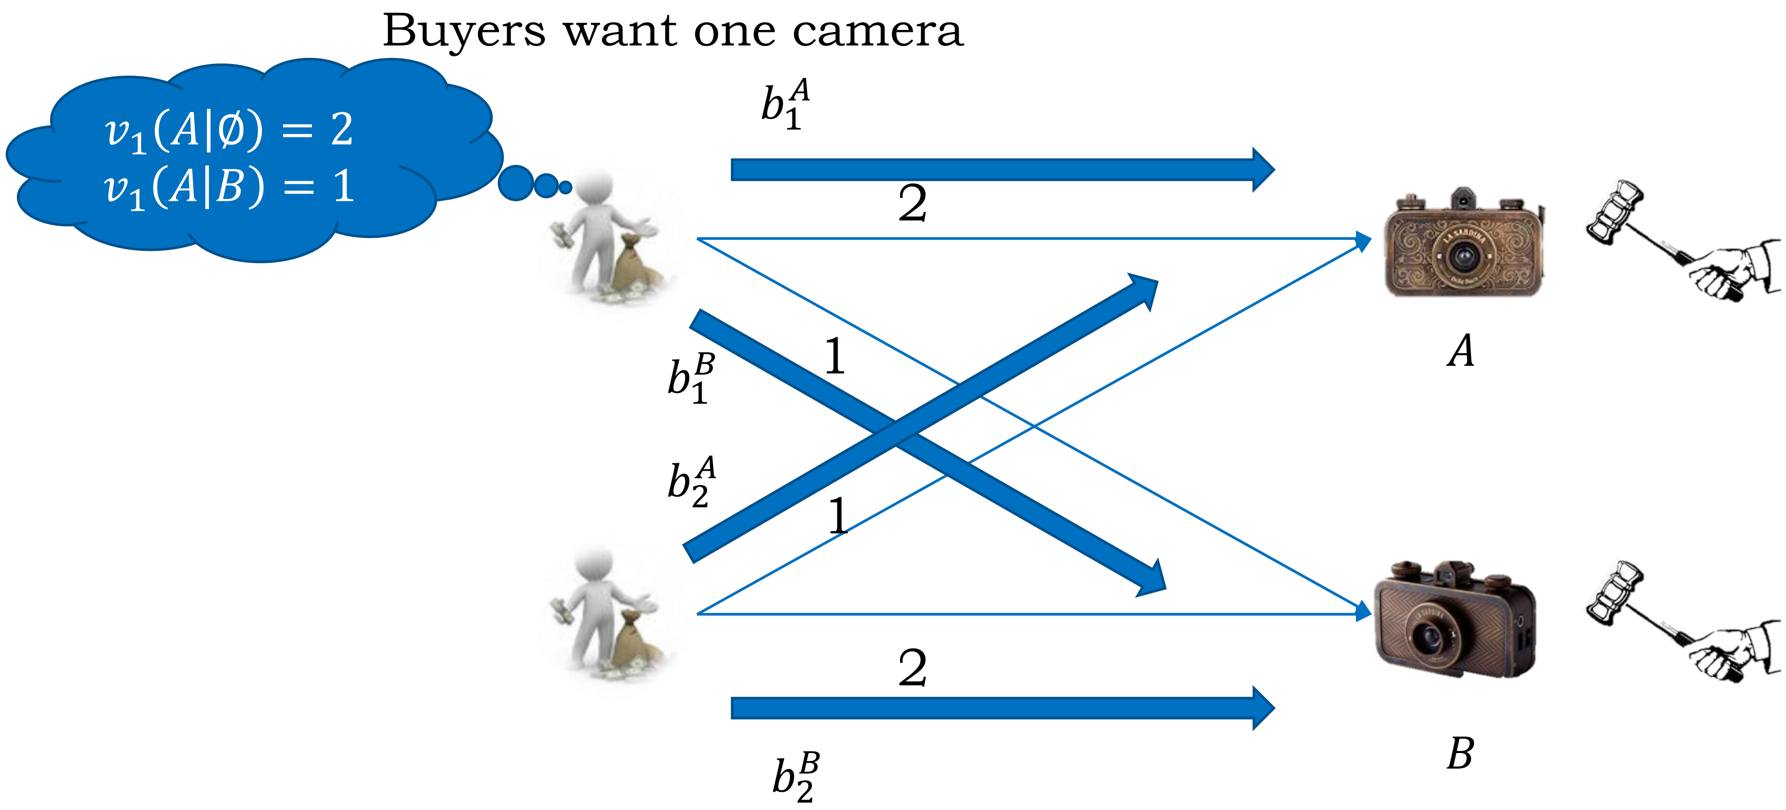

Two second price auctions

Buyers want one camera

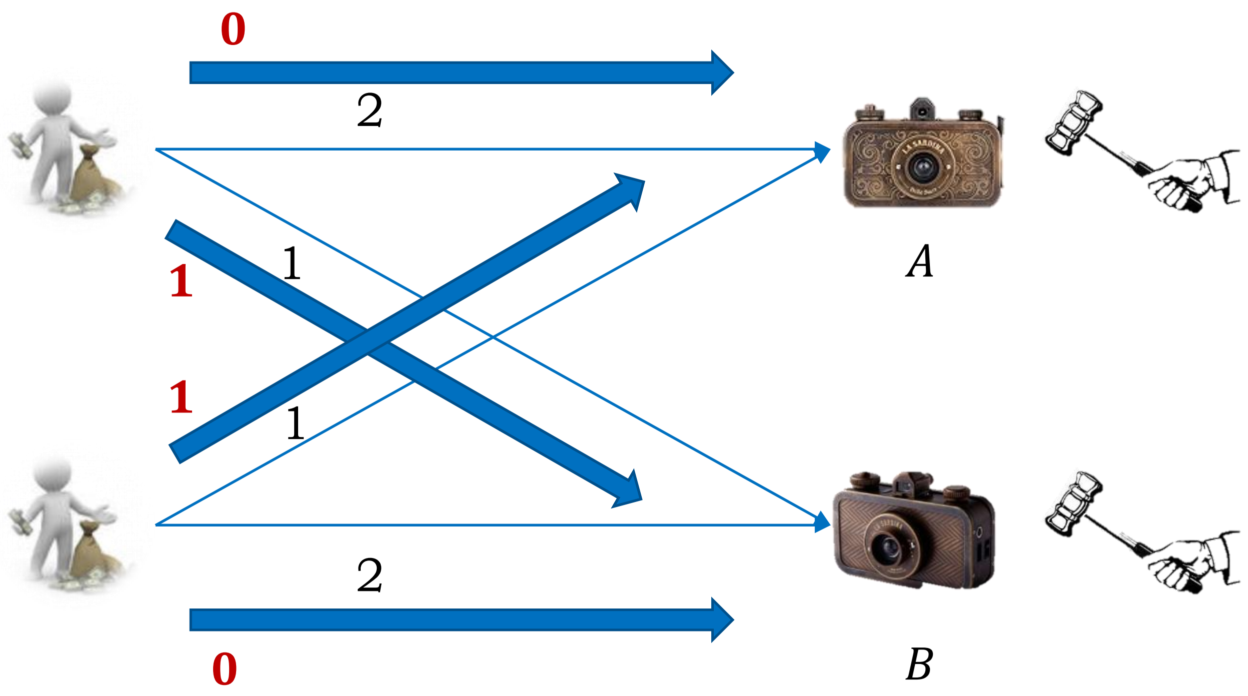

Two second price auctions

Buyers want one camera

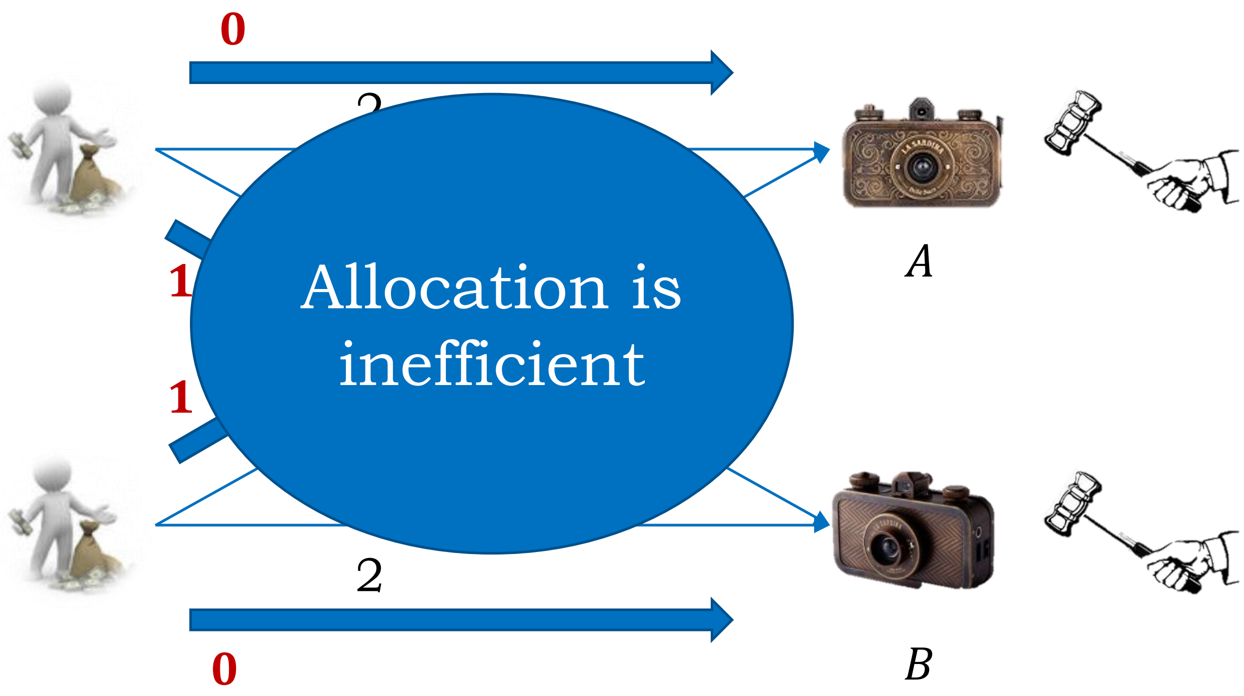

20

Smooth mechanisms

A framework for robust and composable efficiency guarantees

Efficiency at equilibrium

- Truthfulness doesn't compose
- No coordinator to run a centralized global truthful mechanism
- Centralized mechanism too complex or costly to implement

Simple Example: First-Price Auction

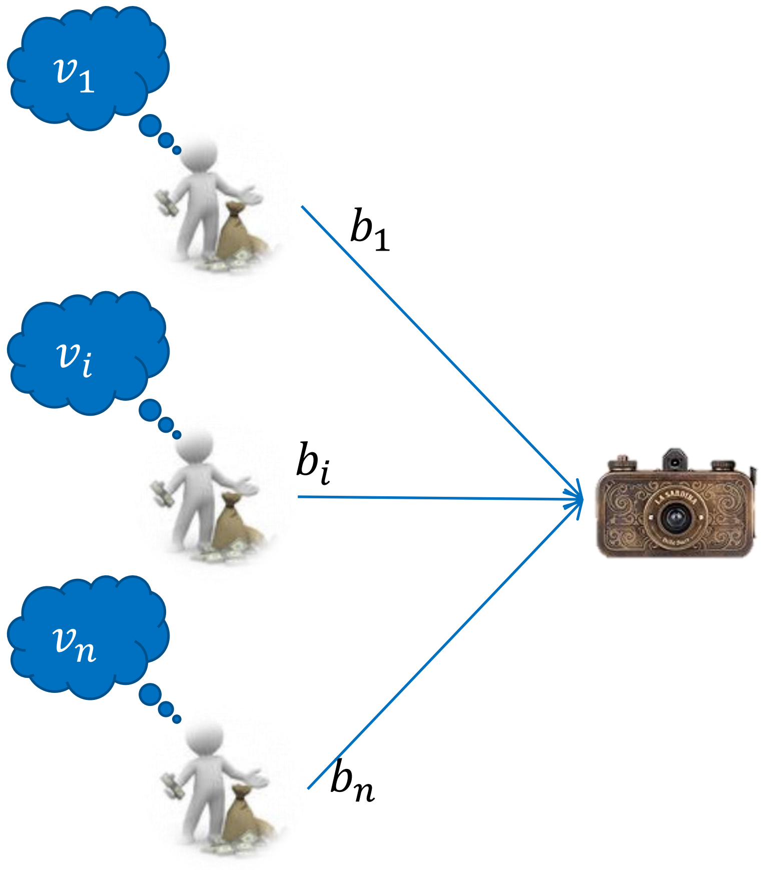

- Utility = Value-Payment:

$$u_i(\mathbf{b}) = (v_i - b_i) \cdot x_i(\mathbf{b})$$

- Efficiency= Welfare

$$SW(\mathbf{b}) = \sum_i v_i \cdot x_i(\mathbf{b})$$

Classic Economics Approach

1. Characterize equilibrium
2. Analyze equilibrium properties

Classic economics approach

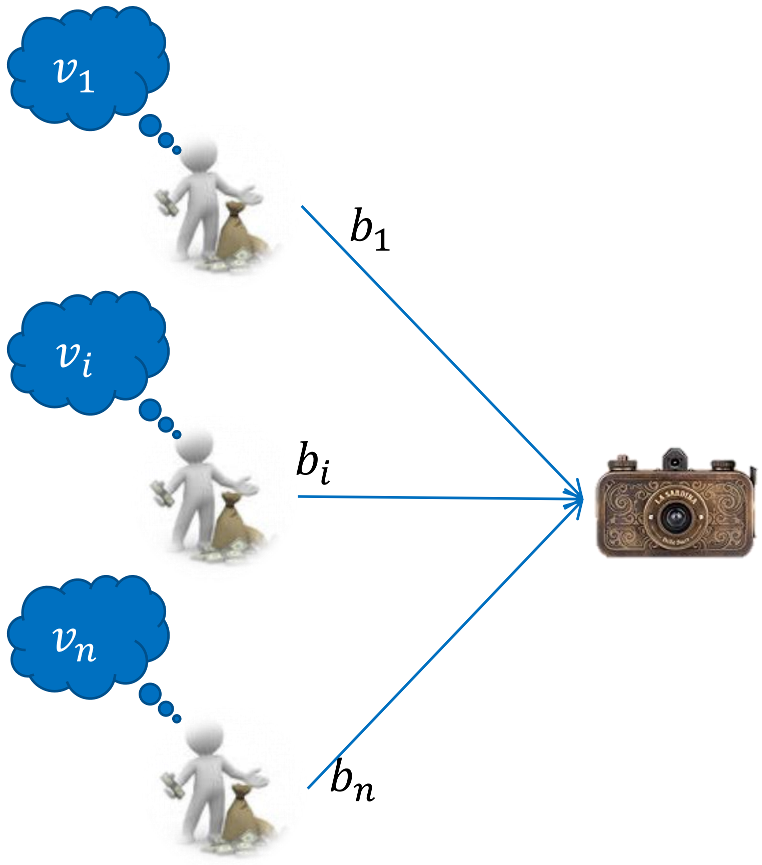

- Characterize equilibrium strategy: b(v)
- Analyze equilibrium properties

Example. $v_i \sim U[0,1]$ then $b(v_i) = \frac{v_i}{2}$

$$u_i(b) = (v_i - b)\Pr[win] = (\mathrm{v_i} - \mathrm{b})2\mathrm{b}$$

Set derivative w.r.t. b equal to 0: $b = \frac{v_i}{2}$

Marvelous theory! Revenue equivalence, Myerson's BNE characterization etc.

One step beyond?

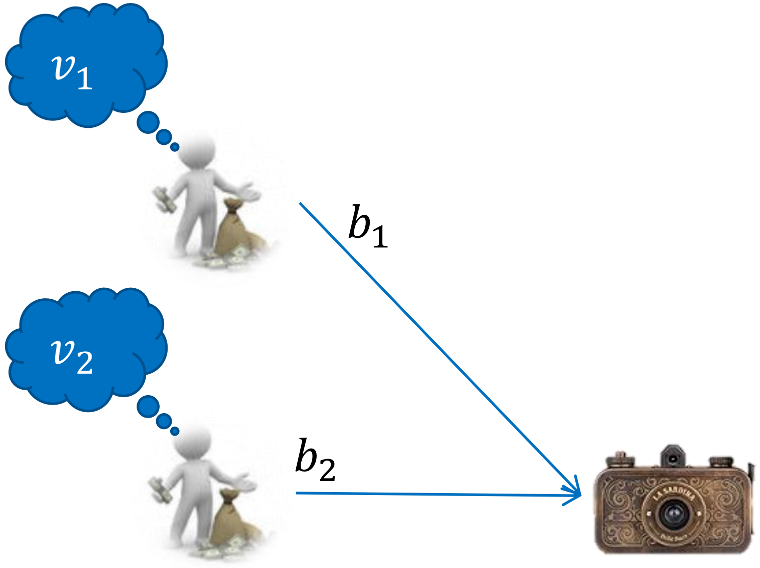

- Characterize equilibrium strategy: b(v)
- Analyze equilibrium properties

Not scalable…

One step beyond?

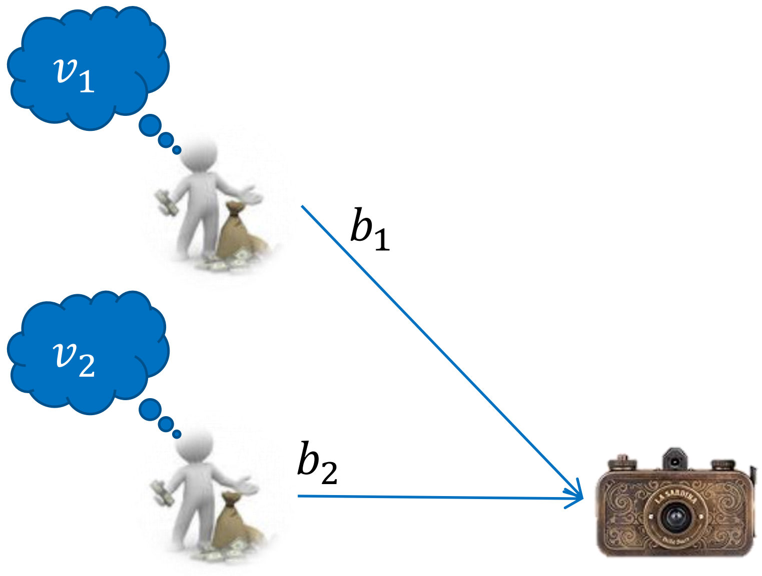

- ~~Characterize equilibrium strategy: b(v)~~
- Analyze equilibrium properties

Not scalable…

One step beyond?

- ~~Characterize equilibrium strategy: b(v)~~
- Analyze equilibrium properties

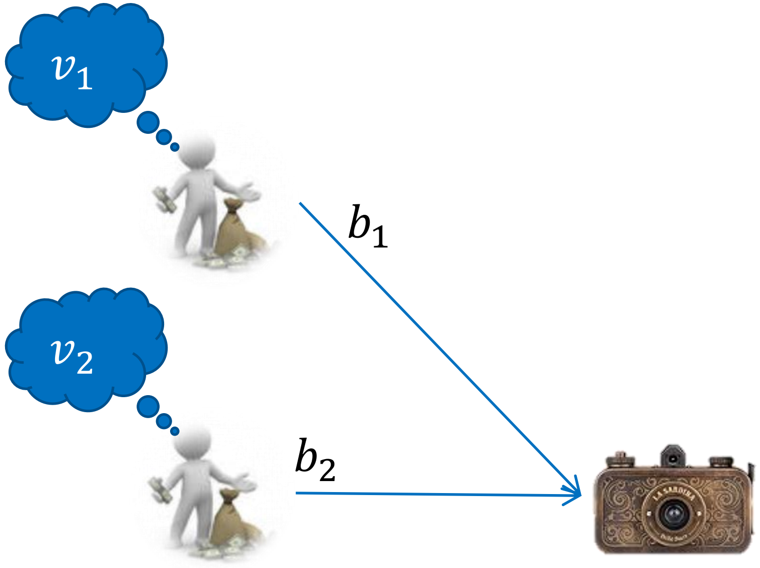

Not scalable…

The Price of Anarchy Approach

Pure Nash Equilibrium and Complete Information

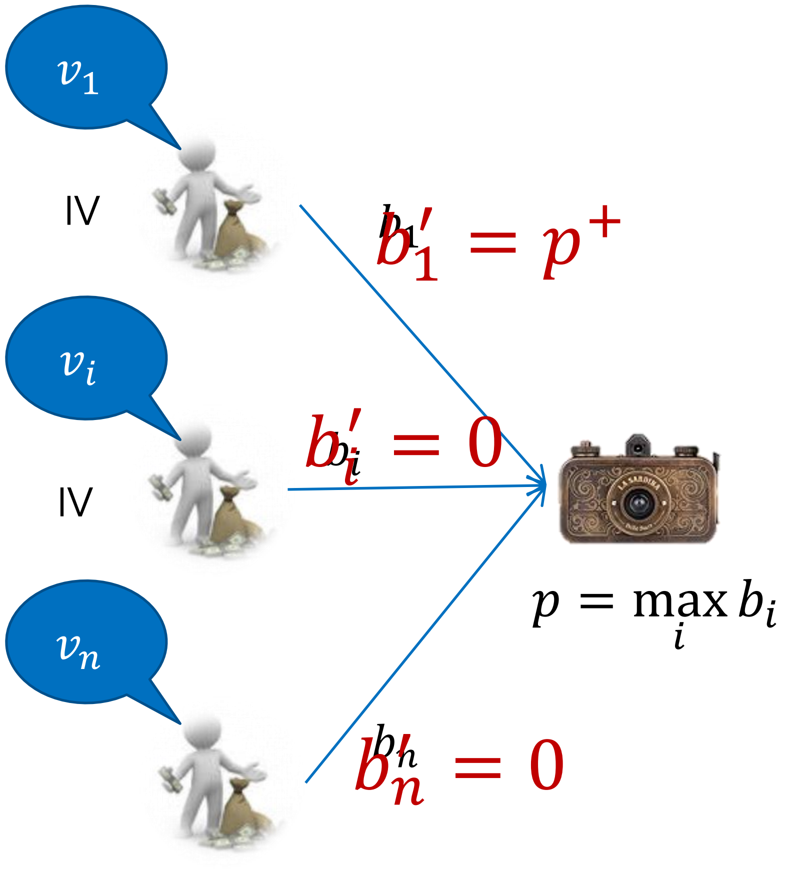

- Pure Nash Equilibrium: b_i maximizes utility

$$u_i(\mathbf{b}) \geq u_i(b_i', \mathbf{b_{-i}})$$

Theorem. Any PNE is efficient.

Proof. Highest value player can deviate to p^+

$$u_1(b) \geq u_1(p^+, \mathbf{b_{-i}}) = v_1 - p^+$$

$$u_i(b) \geq u_i(0, b_{-i}) = 0$$

$$\sum_i u_i(\mathbf{b}) \geq \sum_i u_i(b_i', \mathbf{b_{-i}}) = v_1 - p$$

$$\sum_i v_i x_i(b) - \cancel{p} \geq v_1 - \cancel{p}$$

Robust solution concepts

- Pure Nash of Complete Information is very brittle
  - Pure Nash might not always exist
  - Game might be played repeatedly, with players using learning algorithms (correlated behavior)
  - Players might not know other valuations
  - Players might have probabilistic beliefs about values of opponents

Learning outcomes

$$\begin{array}{cccccc} b_1^1 & b_1^2 & b_1^3 & b_1^4 & b_1^t & b_1^T \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ b_i^1 & b_i^2 & b_i^3 & b_i^4 & b_i^t & b_i^T \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ b_n^1 & b_n^2 & b_n^3 & b_n^4 & b_n^t & b_n^T \end{array}$$

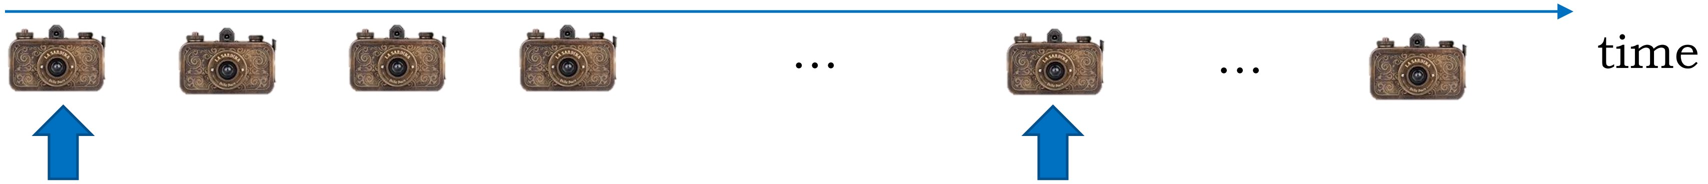

Auction A^1 on $(b_1^1, \dots, b_i^1, \dots, b_n^1)$

Auction A^t on $(b_1^t, \dots, b_i^t, \dots, b_n^t)$

Vanishingly small regret for any fixed strategy x:

$$\sum_{t=1}^{T} u_i(b^t) \geq \sum_{t=1}^{T} u_i\left(x, b_{-i}^t\right) - o(T)$$

Many simple rules: MWU (Hedge), Regret Matching etc.

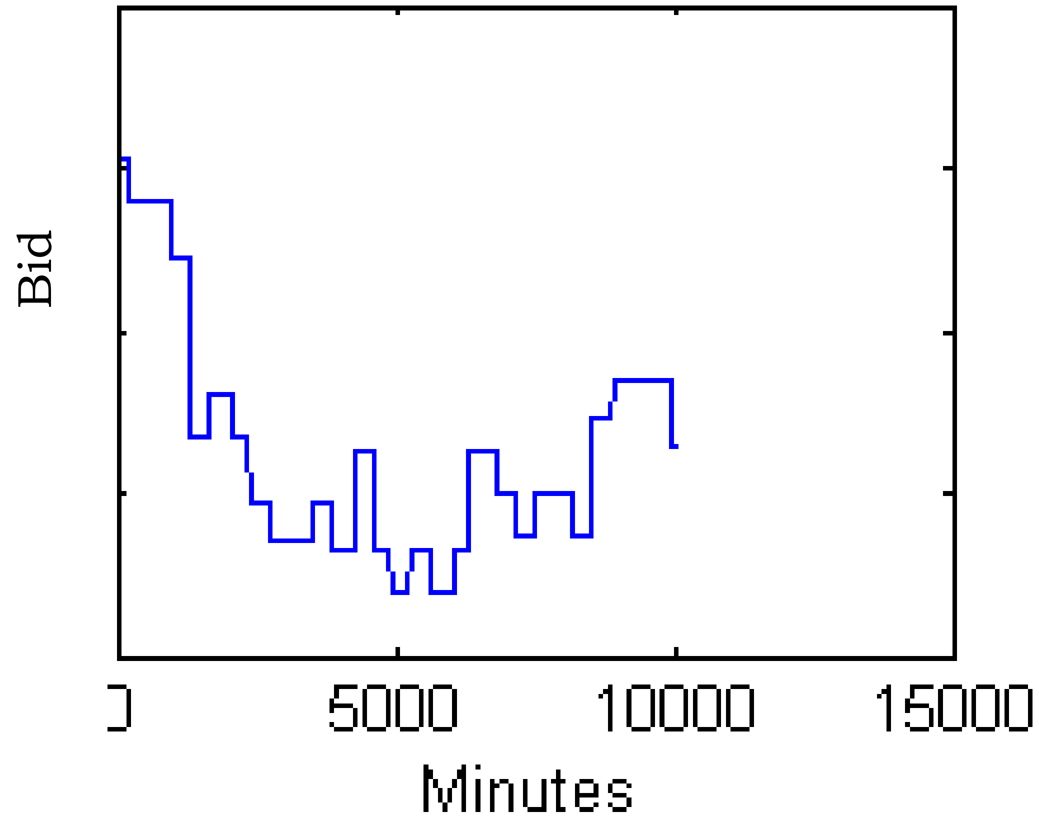
Bid
0
5000
10000
15000
Minutes

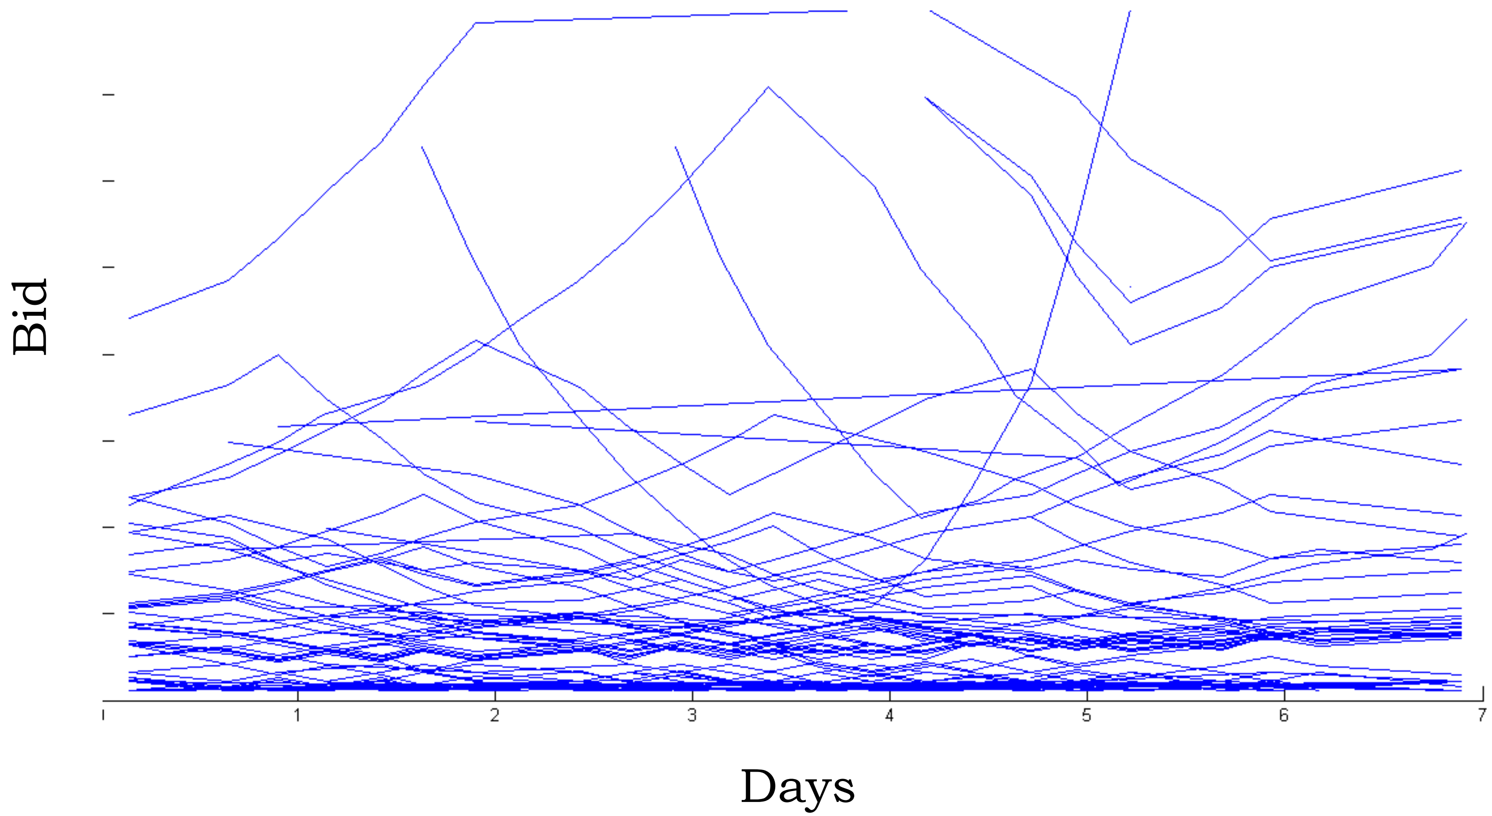

Bid
1
2
3
4
5
6
7
Days

Bayesian beliefs

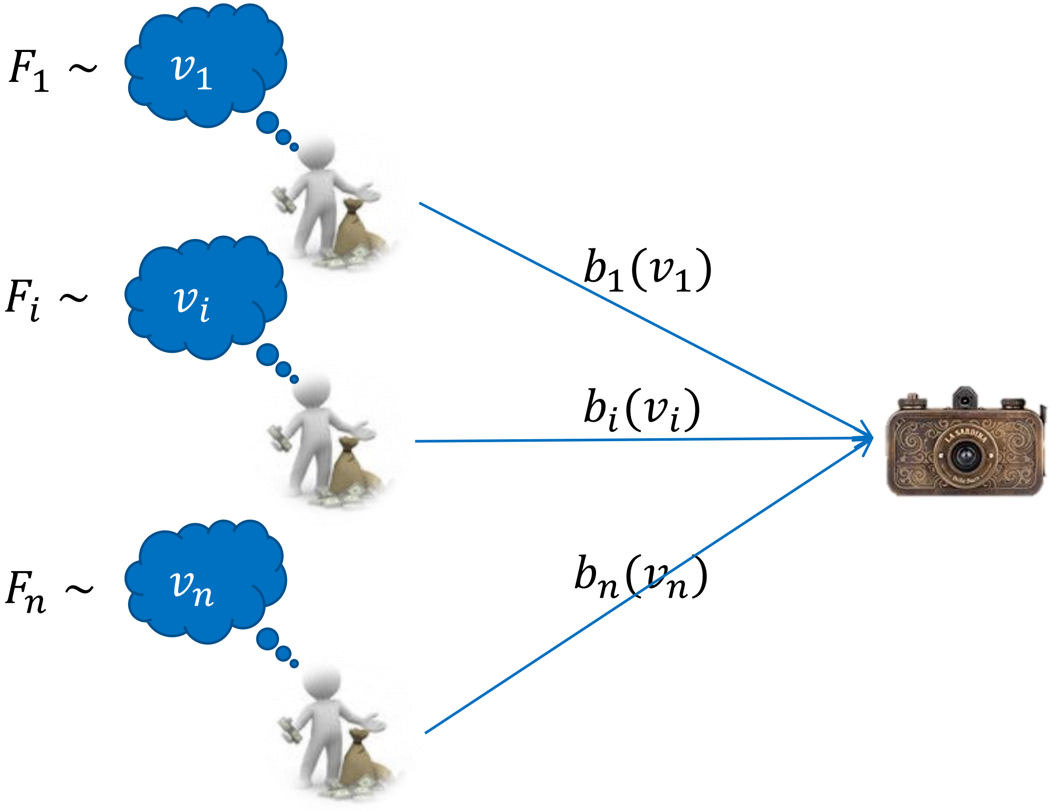

Bayes-Nash Equilibrium:

- Mapping from values to bids
- Maximize utility in expectation

$$E_{v_{-i}}\left[u_i\left(b(v)\right)\right] \geq E_{v_{-i}}\left[u_i\left(b_i', b_{-i}(v_{-i})\right)\right]$$

Expected equilibrium welfare

vs.

Expected ex-post optimal welfare

Direct extensions

- What if conclusions for PNE of complete information directly extended to these more robust concepts
- Obviously: full efficiency doesn't carry over
- Possible, but we need to restrict the type of analysis

Problem in previous PNE proof

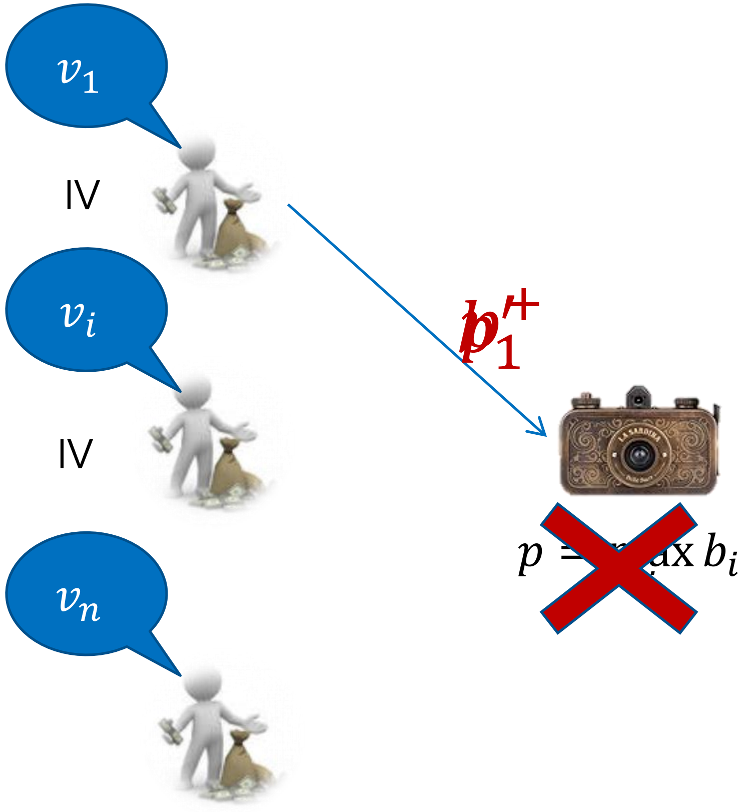

- Recall. PNE is efficient because highest value player doesn't want to deviate to p^+
- Challenge. Don't know p or $\mathbf{v_{-i}}$ in incomplete information
- Idea. Price oblivious deviation analysis
- Restrict deviation to not depend on p

Price-oblivious deviations

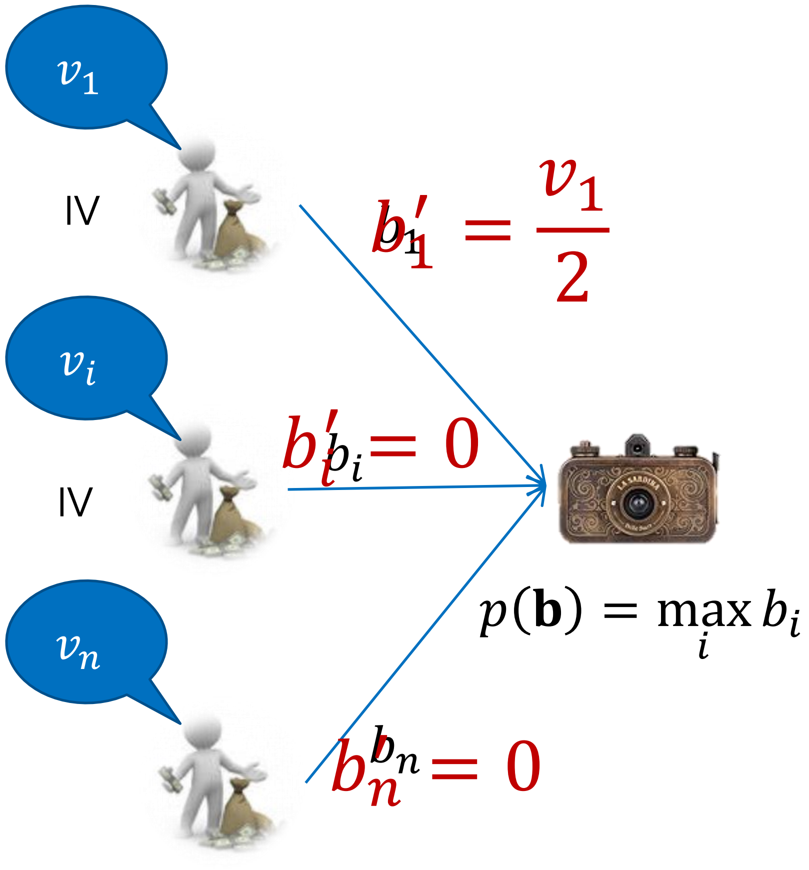

Player 1 can deviate to $b_1' = \frac{v_1}{2}$

- Either $p(\mathbf{b}) \geq \frac{v_1}{2}$
- Or $u_1\left(\frac{v_1}{2}, \mathbf{b_{-1}}\right) = \frac{v_1}{2}$
- In any case:

$$u_1\left(\frac{v_1}{2}, \mathbf{b_{-1}}\right) + p(\mathbf{b}) \geq \frac{v_1}{2}$$

- Others can deviated to $b_i' = 0$:

$$u_i(0, b_{-i}) \geq 0$$

Price-oblivious deviations

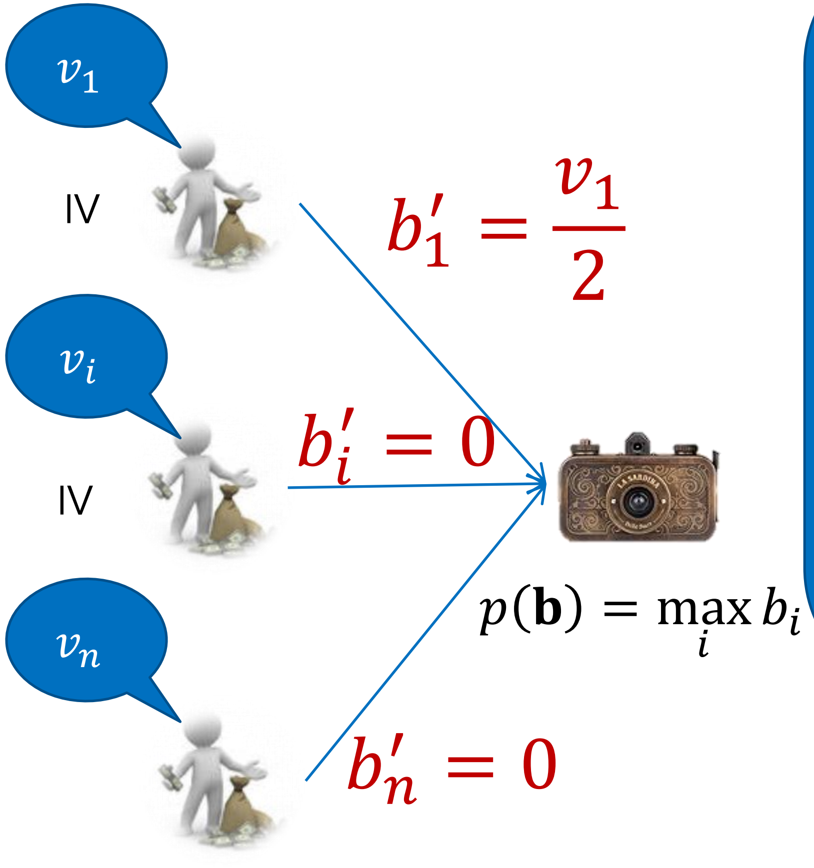

This guarantee extends to learning outcomes and to Bayesian beliefs.

$$SW(\mathbf{b}) \geq \frac{1}{2} OPT(\mathbf{v})$$

Extension to learning outcomes

b_1^1	b_1^2	b_1^3	b_1^4	b_1^t	b_1^T
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_i^1	b_i^2	b_i^3	b_i^4	b_i^t	b_i^T
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
b_n^1	b_n^2	b_n^3	b_n^4	b_n^t	b_n^T

Vanishingly small regret for fixed strategy b_i':

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{i} u_i(b^t) \geq \frac{1}{T}\sum_{t=1}^{T}\sum_{i} u_i(b_i', \mathbf{b}_{-\mathbf{i}}^{\mathbf{t}}) - o(1) \geq \frac{1}{T}\sum_{t=1}^{T}\left(\frac{v_1}{2} - p^t\right) - o(1)$$

$$\frac{1}{T}\sum_{t=1}^{T} SW(b^t) \geq \frac{1}{2} OPT(\mathbf{v}) - o(1)$$

Bayesian Beliefs

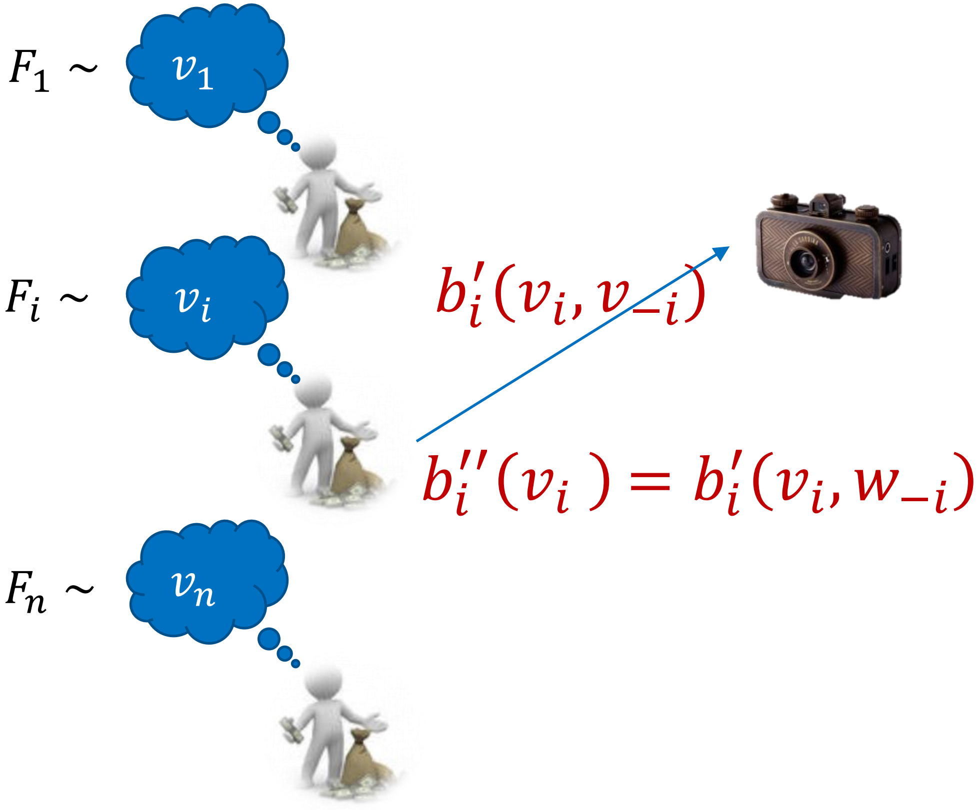

- Deviation depends on opponent values
- Need to construct feasible BNE deviations
- Each player random samples the others values and deviates as if that was the true values of his opponents
- Above works, due to independence of value distributions

Ren
dev

Core Property

Exists b_i' don't depend on current $\mathbf{b}$

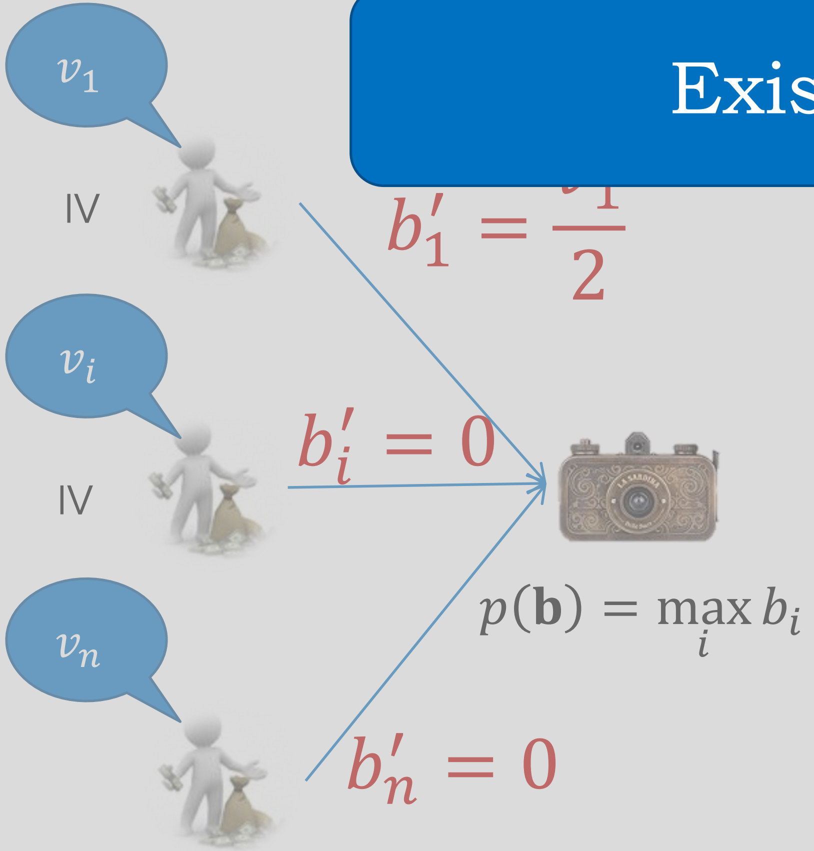

$$\sum_i u_i(\mathbf{b}) \geq \sum_i u_i(b_i', \mathbf{b_{-i}}) \geq \frac{1}{2} OPT(\mathbf{v}) - REV(b)$$

$$\sum_i u_i(\mathbf{b}) \geq \frac{1}{2} v_1 - p(\mathbf{b})$$

$$SW(\mathbf{b}) \geq \frac{1}{2} OPT(\mathbf{v})$$

General mechanisms

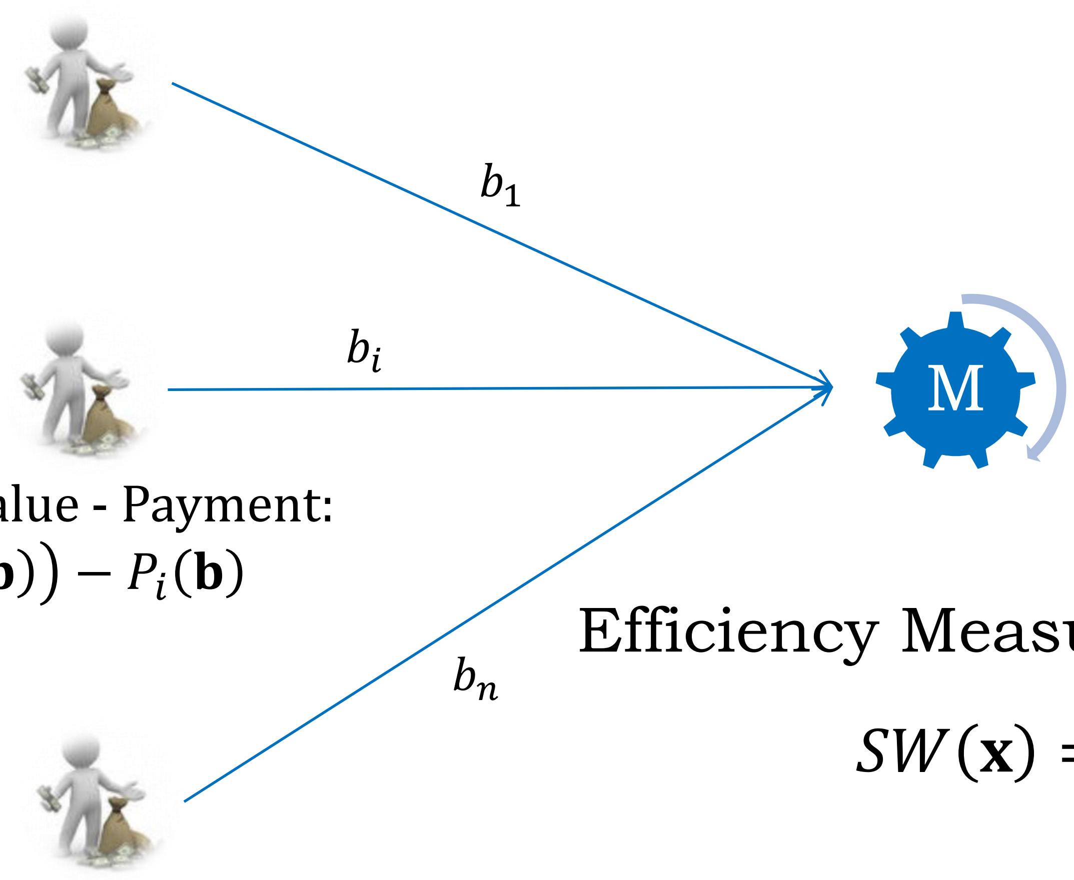

$$X(\mathbf{b}) = \big(X_1(\mathbf{b}), \ldots, X_n(\mathbf{b})\big)$$
$$P(\mathbf{b}) = \big(P_1(\mathbf{b}), \ldots, P_n(\mathbf{b})\big)$$

Utility = Value - Payment:

$$v_i\big(X_i(\mathbf{b})\big) - P_i(\mathbf{b})$$

Efficiency Measure: Social Welfare

$$SW(\mathbf{x}) = \sum_i v_i(x_i)$$

General mechanisms

Utility = Value - Payment:

$v_i\big(X_i(\mathbf{b})\big) - P_i(\mathbf{b})$

Combinatorial Auctions

General mechanisms

b_1

b_i

Utility = Value - Payment:
$v_i(X_i(\mathbf{b})) - P_i(\mathbf{b})$

b_n

Public Projects

General mechanisms

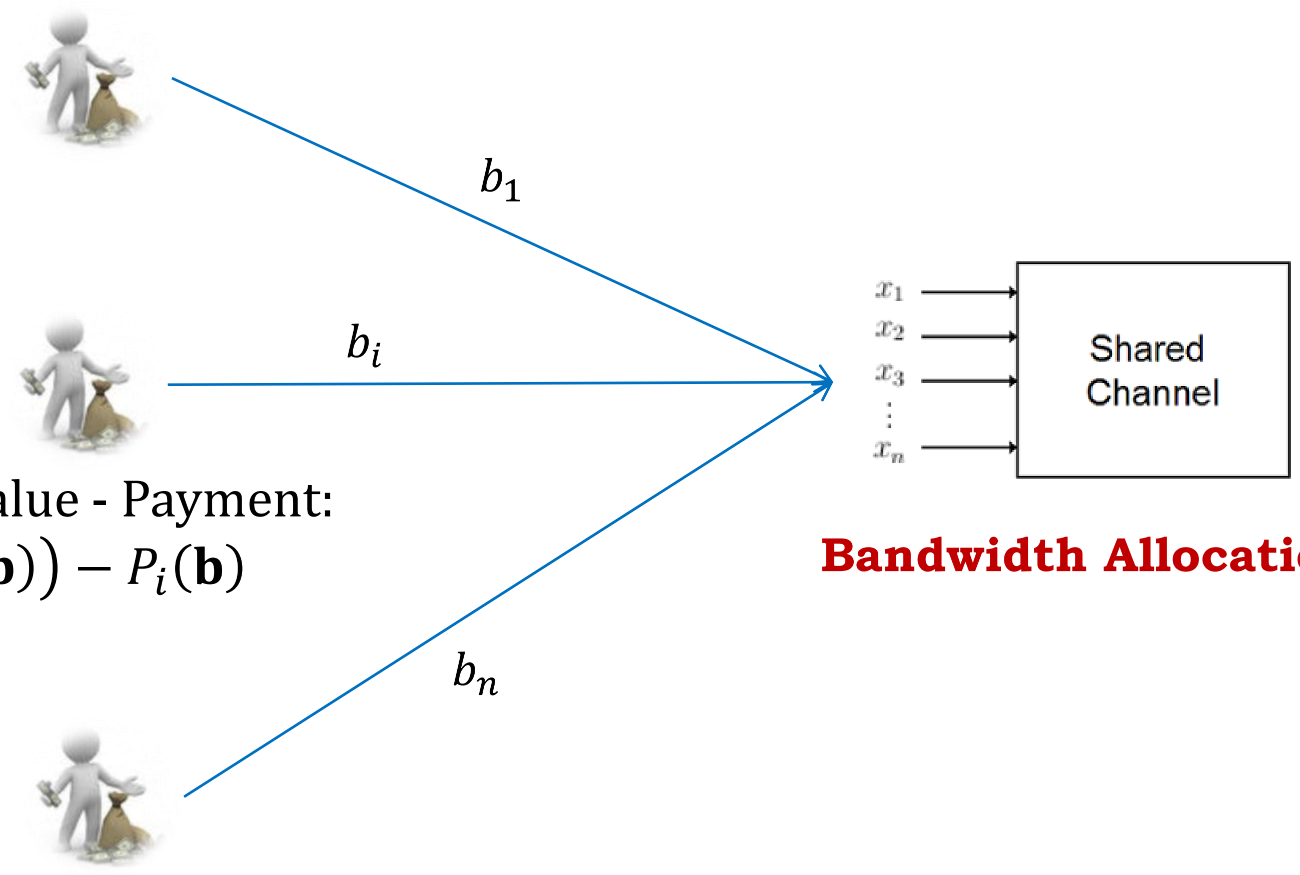

Utility = Value - Payment:

$v_i(X_i(\mathbf{b})) - P_i(\mathbf{b})$

Bandwidth Allocation

General mechanisms

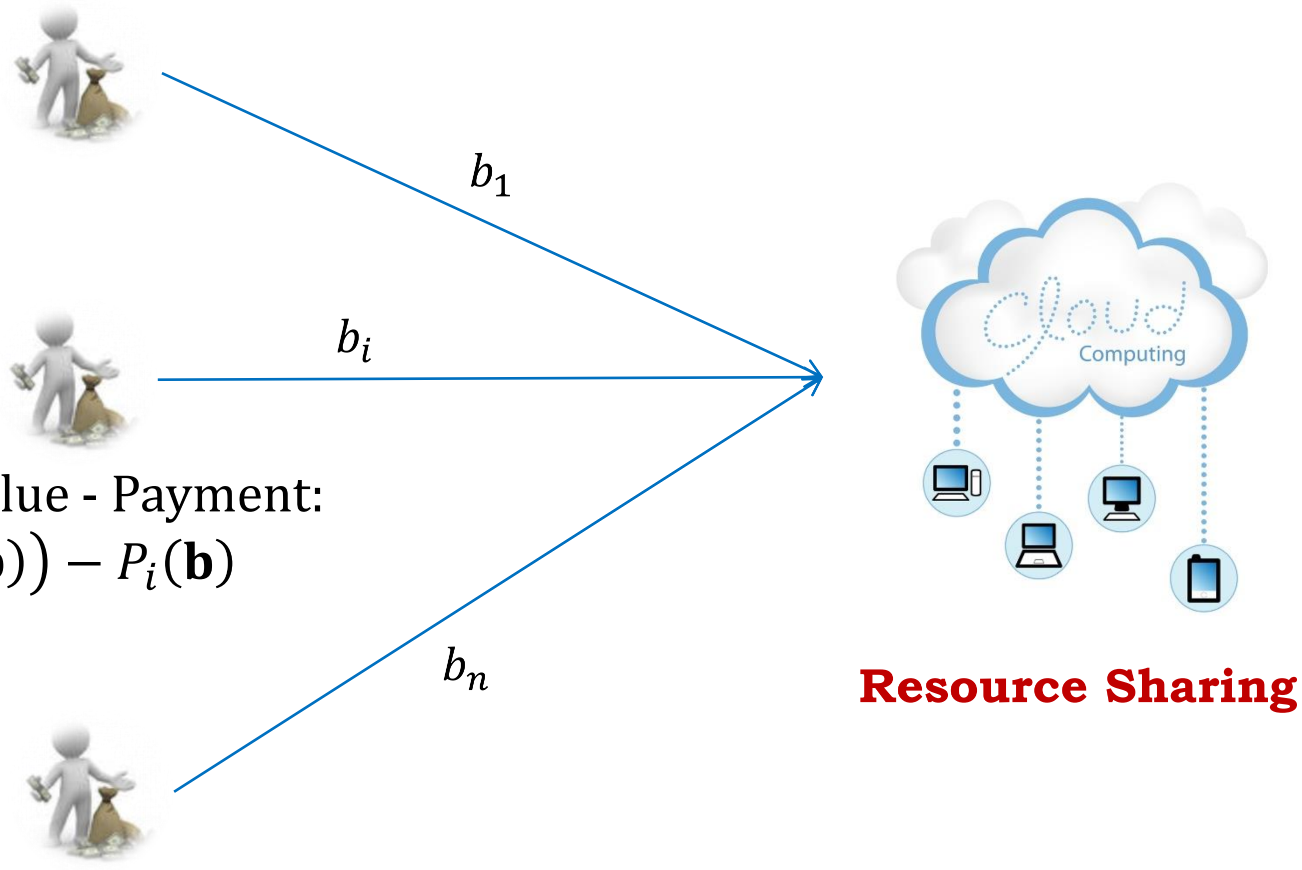

Utility = Value - Payment:
$v_i(X_i(\mathbf{b})) - P_i(\mathbf{b})$

Resource Sharing

Definition (S.-Tardos'13)

(λ, μ) –Smooth Mechanism

Exist b_i' that don't depend on current $\mathbf{b}$

For any $\mathbf{b}$

$$\sum_i u_i(b_i', \mathbf{b_{-i}}) \geq \boldsymbol{\lambda} \cdot OPT(\mathbf{v}) - \boldsymbol{\mu} \cdot REV(\mathbf{b})$$

Closely related to smooth games [Roughgarden STOC'09], giving an intuitive market interpretation of smoothness

Robust Efficiency Guarantees

Theorem (S-Tardos'13) Mechanism is (λ, μ)-smooth, then every Nash Equilibrium achieves at least $\frac{\lambda}{\max\{1, \mu\}}$ of OPT.

- Extends no-regret learning outcomes of repeated game
- Extends to **Bayesian Setting**, assuming **independent value distributions** and even to no-regret under incomplete information.
  - Extending Roughgarden EC'12 and **S.**'12 that used stricter universal smoothness property

Simple Mechanisms in the Literature

- **Simultaneous Second Price Single-Item Auctions**
  Christodoulou, Kovacs, Schapira ICALP'08, Bhawalkar, Roughgarden SODA'11
- **Auctions based on Greedy Allocation Algorithms**
  Lucier, Borodin SODA'10
- **AdAuctions (GSP, GFP)**
  Paes-Leme Tardos FOCS'10, Lucier, Paes-Leme + CKKK EC'11
- **Simultaneous First Price Auctions Single-Item Auctions**
  Bikhchandani GEB'96, Hassidim, Kaplan, Mansour, Nisan EC'11, Fu et al. STOC'13
- **Sequential First/Second Price Auctions**
  Paes Leme, S, Tardos SODA'12, S, Tardos EC'12

All above can be thought as smooth mechanisms and some are even compositions of smooth mechanisms.

Applications of Smooth Mechanisms

- First price auction: $\left(1 - \frac{1}{e}, 1\right)$-smooth (Improves Hassidim et al. EC'12)
- First price combinatorial auction based on a a-approximate greedy algorithm is $(1 - e^{-a}, 1)$-smooth (Improves Lucier-Borodin SODA'10)
- Marginal pricing multi-unit auctions is $\left(1 - \frac{1}{e}, 1\right)$-smooth (Improves De Keijzer et al. ESA'13)
- All-pay auction: $\left(\frac{1}{2}, 1\right)$-smooth (New result)
- First price position auction is $\left(\frac{1}{2}, 1\right)$-smooth
  - Extends Paes Leme et al. FOCS'10 to more general valuations)
- Proportional bandwidth allocation mechanism is $\left(\frac{1}{2}, 1\right)$-smooth
  - Extends Johari-Tsitsiklis'05, to incomplete information and learning outcomes

Composition of Mechanisms

A simple example: Simultaneous First-Price Auctions

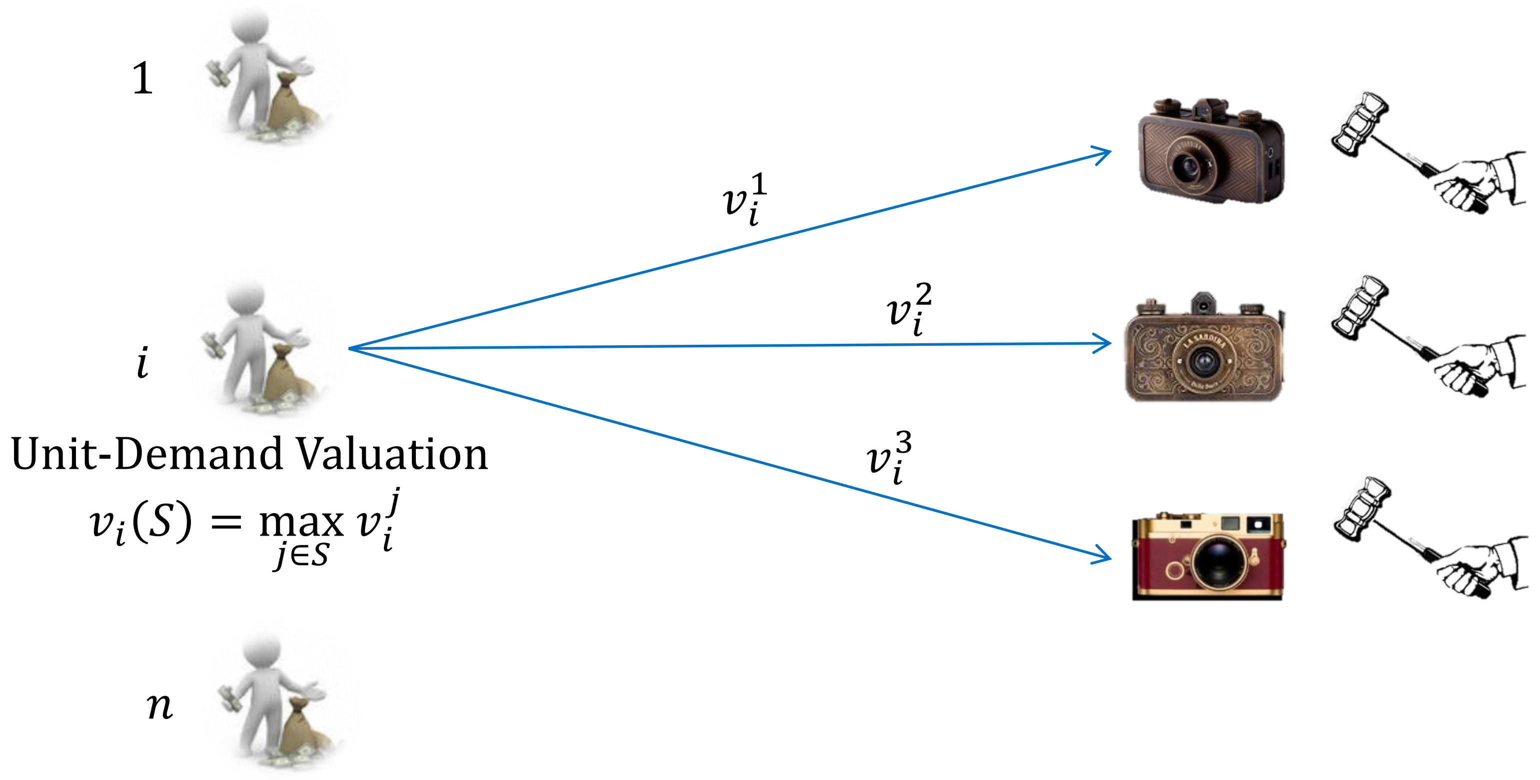

A simple example: Simultaneous First-Price Auctions

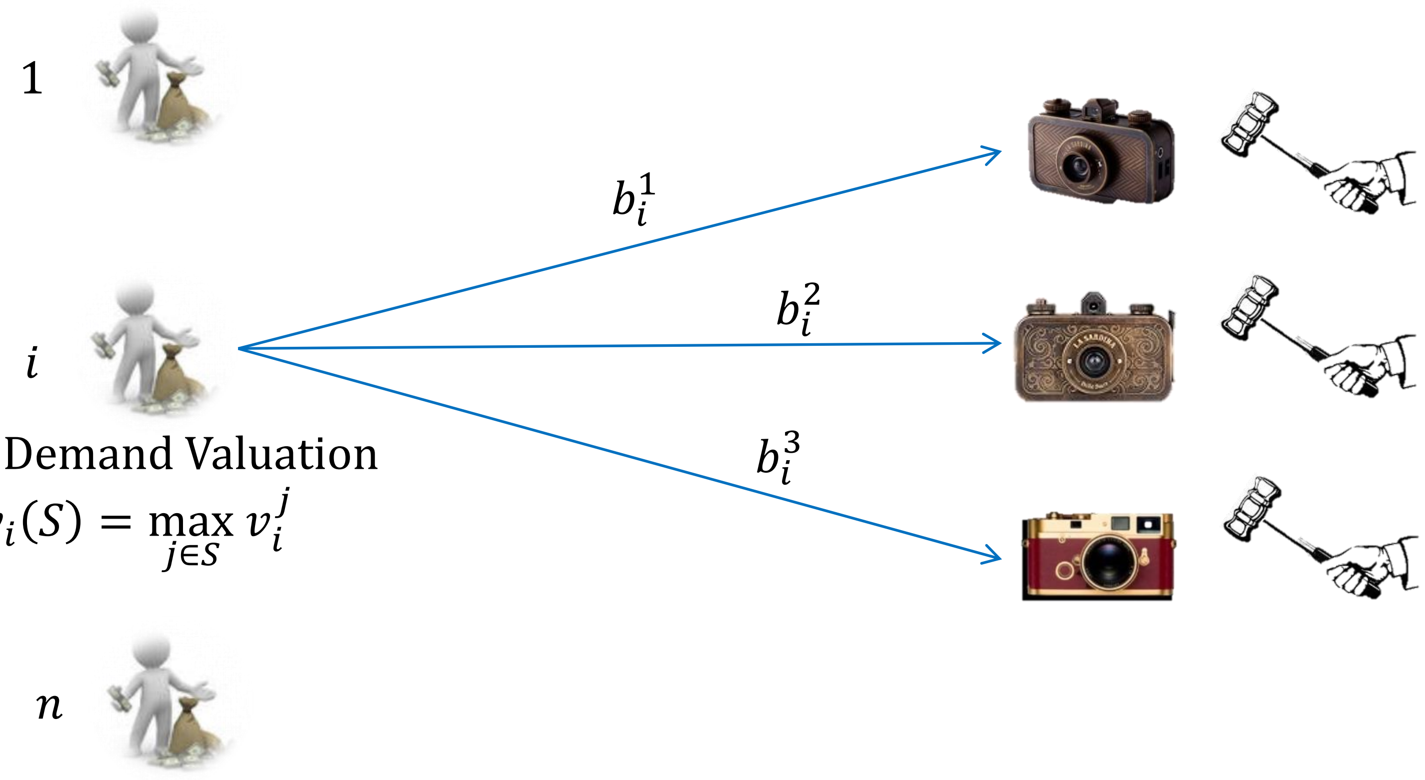

Unit-Demand Valuation

$$v_i(S) = \max_{j \in S} v_i^j$$

Global efficiency guarantees

Can we derive global efficiency guarantees from local $\left(\frac{1}{2},1\right)$ −smoothness of each first price auction?

APPROACH: Prove smoothness of the global mechanism

GOAL: Construct global deviation

IDEA: Pick your item in the optimal allocation and perform the smoothness deviation for your local value v_i^j, i.e. $\frac{v_i^j}{2}$

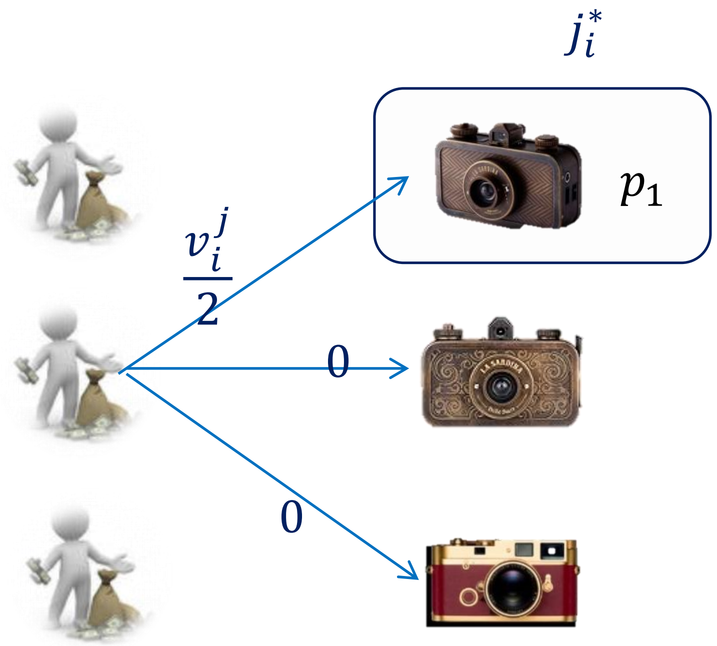

Local to Global Smoothness

Smoothness locally:

$$u_i(b_i', \mathbf{b_{-i}}) \geq \frac{v_i^{j_i^*}}{2} - p_{j_i^*}(\mathbf{b})$$

Summing over players:

$$\sum_i u_i(b_i', \mathbf{b_{-i}}) \geq \frac{1}{2} \cdot OPT(\mathbf{v}) - REV(\mathbf{b})$$

Implying $\left(\frac{1}{2}, 1\right)$ −smoothness property globally.

Composition of Mechanisms

$X^1(\mathbf{b^1})$
$P^1(\mathbf{b^1})$

$X^j(\mathbf{b^j}) = \left(X_1^j(\mathbf{b^j}), \ldots, X_n^j(\mathbf{b^j})\right)$
$P^j(\mathbf{b^j}) = \left(P_1^j(\mathbf{b^j}), \ldots, P_n^j(\mathbf{b^j})\right)$

$X^m(\mathbf{b^m})$
$P^m(\mathbf{b^m})$

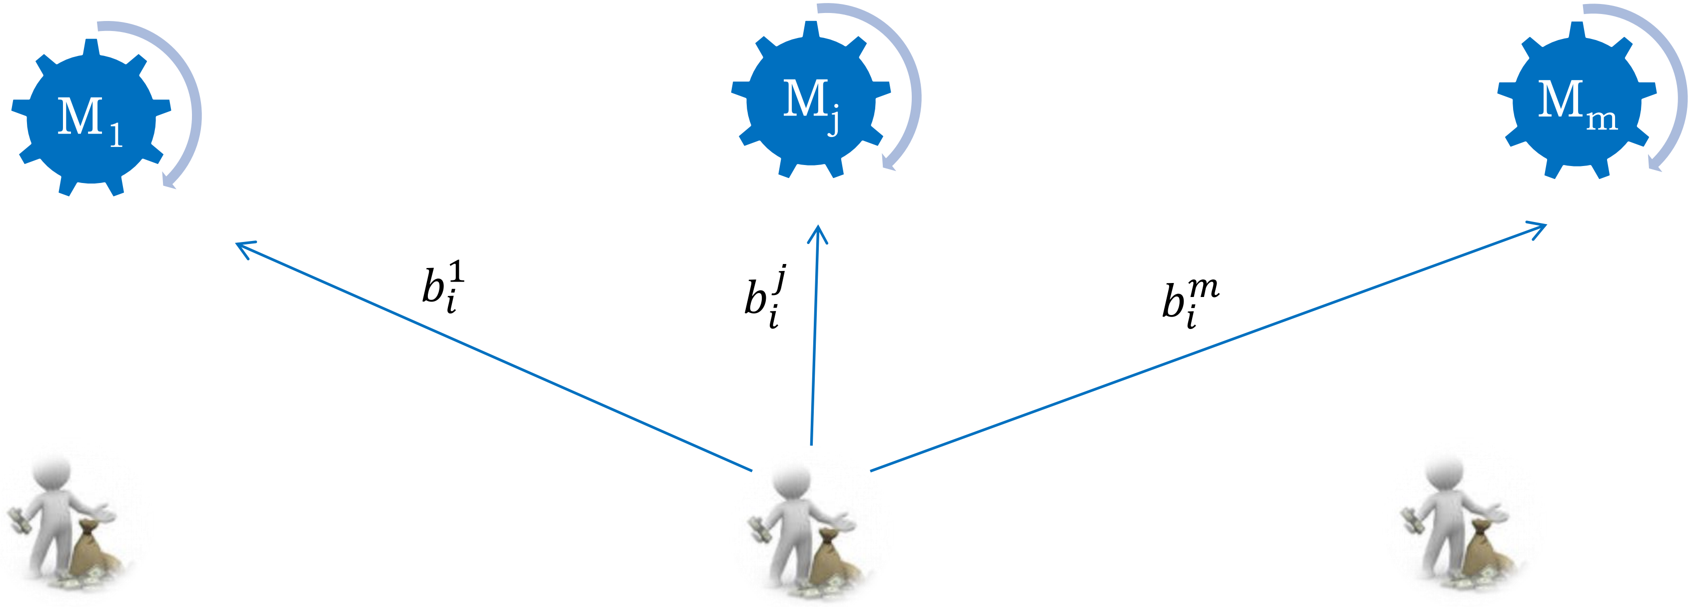

Complex valuation over outcomes
$v_i\left(X_i^1(\mathbf{b^1}), \ldots, X_i^m(\mathbf{b^m})\right)$

Simultaneous Composition

Theorem (S.-Tardos'13) Simultaneous composition of m mechanisms, each (λ, μ)-smooth and players *have no complements* across mechanisms*, then composition is also (λ, μ)-smooth.

No-complements Across Mechanisms

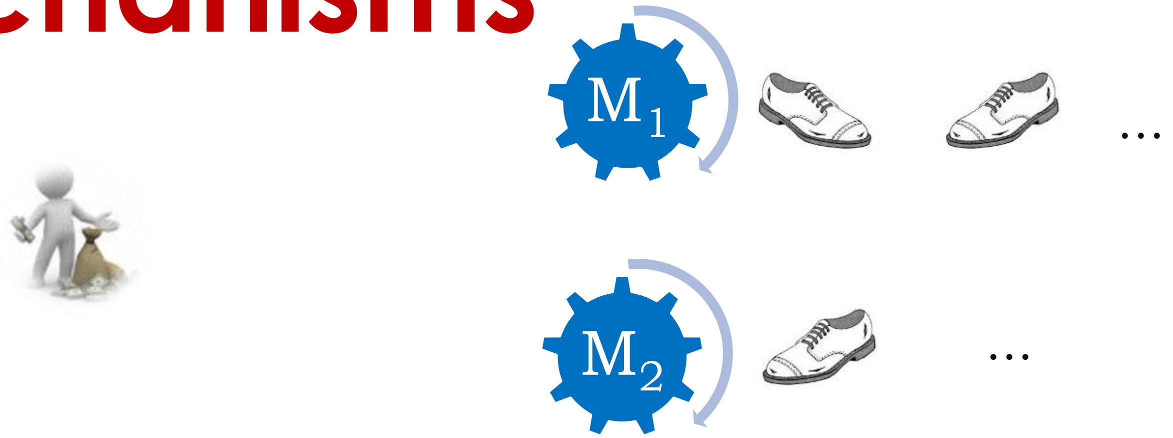

- Marginal value for any allocation from some mechanism can only decrease, as I get non-empty allocations from more mechanisms

- No assumption about allocation structure and valuation within mechanism

Global Efficiency Theorem.

A market composed of (λ, μ)-Smooth Mechanisms achieves $\frac{\lambda}{\max\{1,\mu\}}$ of optimal welfare at no-regret learning outcomes and under incomplete information, when players have no-complement valuations across mechanisms.

Extensions

- Sequential Composition

  Smooth mechanisms compose sequentially when values are unit-demand*.

  Tight via: Feldman, Lucier, **S.** “Limits of Efficiency in Sequential Auctions”

- Hard Budget Constraints on Payments

  Same efficiency guarantees with respect to new welfare benchmark:

  Optimal welfare achievable after capping a player’s value by his budget

- Limited complementarities
  - Global efficiency degrades smoothly with size of complementarities
  - Feige, Feldman, Immorlica, Izsak, Lucier, **S.**, “A Unifying Hierarchy of Valuations with Complements and Substitues”

Open problems – Recent results

- Revenue of non-truthful mechanisms via price of anarchy in multi-dimensional settings
  - "Price of anarchy for auction revenue": Hartline, Hoy, Taggart
- Other models of non-fully rational behavior: level-k, fictitious play
  - "Level-0 Meta-Models for Predicting Human Behavior in Games": J. Wright, K. Leyton-Brown
- Simple auctions with simple strategies: good mechanisms with small strategy spaces (single knob to turn, simple to optimize over)
  - "Utility target mechanisms": Hoy, Jain, Wilkens
  - "Simple auctions with simple strategies": Devanur, Morgenstern, **S.**, Weinberg
- Algorithmic characterization of smoothness in multi-dimensional environments (similar to cyclic monotonicity)
- Uncertainty about own valuation, information asymmetry
  - "Auctions, Adverse Selection, and Internet Display Advertising", Arnosti, Beck, Milgrom
- Coalitional dynamics – analogues of no-regret dynamics with good welfare properties
  - "Strong Price of Anarchy and Coalitional Dynamics": Bachrach, **S.**, Tardos, Vojnovic

In brief

- Many simple mechanisms are smooth
- Smooth mechanisms compose well
- Robust efficiency guarantees
- Useful design and analysis tool for efficiency in electronic markets/distributed resource allocation systems

Thank you!