

When Are Equilibria of Simple Auctions Near-Optimal?

Part 1: Positive Results

Tim Roughgarden (Stanford)

*Reference: Roughgarden/Syrgkanis/Tardos,
“The Price of Anarchy in Auctions,” Journal of
Artificial Intelligence Research (JAIR), 2017*

Multi-Item Auctions

- n bidders (e.g. telecoms), m items (e.g. licenses)
- bidder i has private nonnegative *valuation* $v_i(S)$ for each subset S of items [$\approx 2^m$ parameters!]
- bidder i wants to maximize $v_i(S_i) - \text{payment}$
 - .
- *social welfare* of allocation S_1, S_2, \dots, S_n : $\sum_i v_i(S_i)$
 - objective is to allocate items to maximize this

Preamble

A Simple Auction Format

First cut: [McAfee, Milgrom-Wilson 93] *simultaneous ascending auctions* (one auction per item).

- usually works decently, but:

Issue #1: *demand reduction.*

- bidder buys fewer items to get a cheaper price

Issue #2: *exposure problem.*

- example: 2 items; bidder #1 has value 6 for *both* items, bidder #2 wants one item, value = 5

Inefficiency in Auctions

(P. Cramton, “The Efficiency of FCC Spectrum Auctions,” 1998)

“The setting of spectrum auctions is too complex to guarantee full efficiency.”

“Direct evidence of demand reduction was seen in the nationwide narrowband auction. The largest bidder, PageNet, reduced its demand from three of the large licenses to two, at a point when prices were still well below its marginal valuation for the third unit. PageNet felt that, if it continued to demand a third license, it would drive up the prices on all the others to disadvantageously high levels.”

“Nonetheless, an examination of the bidding suggests that these problems, although present, probably did not lead to large inefficiencies.”

Inefficiency in Auctions

“[T]he measured efficiency of the simultaneous ascending auction falls off markedly as complementarities increase, but the efficiency of the package auction is largely unaffected by complementarity.” (L. M. Ausubel and P. R. Milgrom, “Ascending Auctions with Package Bidding”, 2002.)

Practical Rules of Thumb

Folklore belief #1: without strong complements, simple auctions work pretty well.

- *loss in outcome quality appears small*
- *demand reduction exists, but not a dealbreaker*

Folklore belief #2: with strong complements, simple auctions aren't good enough.

- *loss in outcome quality could be big*
- *exposure problem exists, and is a dealbreaker*

A Representative Result

Example Theorem: [Syrkanis/Tardos 13] (improving [Hassidim/Kaplan/Nisan/Mansour 11]) Suppose m items are sold simultaneously via first-price single-item auctions:

- for every product distribution over submodular bidder valuations (independent, not necessarily identical), and
- for every (mixed) Bayes-Nash equilibrium, expected welfare of the equilibrium is within 63% of the maximum possible.

(submodular = decreasing marginal values)

First-Price Auctions

First-Price Auctions

- one item, n players
- winner = highest bidder, price = highest bid
- player i has private valuation v_i , quasi-linear utility

Common Prior Assumption: valuations drawn from a distribution known to all players (independent, or not).

- *strategy*: function from valuations to bids
 - semantics: “if my valuation is v , then I will bid b ”

Bayes-Nash Equilibrium: every player i picks expected utility-maximizing action, given its knowledge (prior, own valuation, others' strategies).

The Role of Symmetry

Bayes-Nash Equilibrium: every player picks expected utility-maximizing action, given its knowledge.

Exercise: with n bidders, valuations drawn i.i.d. from $U[0,1]$, the following is a Bayes-Nash equilibrium: all bidders use the strategy $v_i \rightarrow [(n-1)/n] \cdot v_i$.

- highest-valuation player wins (maximizes welfare)

Exercise: with 2 bidders, valuations from $U[0,1]$ and $U[0,2]$, no Bayes-Nash equilibrium maximizes expected welfare. (Second bidder shades bid more.)

Correlated Valuations

Example: [Syrkanis 14] 3 bidders.

- $v_1 = 1$ (deterministic)
- $v_2=v_3$ with distribution fn $F(v) = 1/(e(1-t))$ on $[0,1-1/e]$

Expected optimal welfare: 1 (give item to bidder 1).

Bayes-Nash equilibrium:

- bidder 1 bids 0 (assume ties broken in favor of #1)
- bidders 2, 3 bid truthfully
- expected welfare = $1-1/e \approx 63\%$

Digression: Multiple Items

Example: [Engelbrecht-Wiggins/Weber 79]

- n bidders, n items
- each bidder has value 1 for each item (only wants one)
- mechanism: each bidder bids on a single item, each item allocated via first-price auction

Optimal welfare: n (give one item to each bidder).

Symmetric mixed Nash equilibrium:

- each bidder picks an item uniformly at random
- bid as in symmetric FPA with random # of bidders
- expected welfare $\rightarrow 1 - 1/e \approx 63\%$ (as $n \rightarrow \infty$)

A Terminological Aside

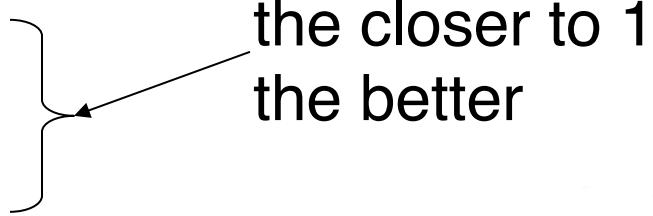
Price of Anarchy: Definition

Definition: [Koutsoupias/Papadimitriou 99]

price of anarchy (POA) of a game (w.r.t. some objective function, equilibrium concept):

$$\frac{\text{equilibrium objective fn value}}{\text{optimal obj fn value}}$$

the closer to 1
the better



- if multiple equilibria, defined by the worst (farthest-from-optimal) one

The POA Goes Viral

Example domains: scheduling, routing, facility location, bandwidth allocation, network formation, network cascades, contention resolution, coordination games, firm competition, auctions, ...

The Price of Anarchy of Health Care

Dr. Vincent Knight, Professor Paul Harper

Cardiff University, Cardiff, UK

knightva@cf.ac.uk

The price of anarchy in basketball

Brian Skinner

School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455

Back to First-Price Auctions

Equilibrium Welfare Guarantees

Proof plan:

- fix an arbitrary equilibrium $\mathbf{b} = (b_1, \dots, b_n)$
- choose i 's and b^*_i 's to derive inequalities of form
 $u_i(\mathbf{b}) \geq u_i(b^*_i, \mathbf{b}_{-i})$ [b^*_i 's = *baseline strategies*]
 - guideline: \mathbf{b}^* typically induces welfare-maximizing outcome
- compile inequalities into one of the form
welfare(\mathbf{b}) $\geq \lambda \cdot (\text{max-possible welfare})$

For best results: choose baselines b^*_i 's
independently of \mathbf{b} .

Guarantees for PNE

Assume: for suitable choice of \mathbf{b}^* , for every \mathbf{b} ,

$$\sum_i u_i(b_i^*, \mathbf{b}_{-i}) \geq \lambda \cdot [\text{OPT Welfare}] - \text{Rev}(\mathbf{b}).$$

Claim: POA of pure Nash equilibria is $\geq \lambda$.

Proof: Let \mathbf{b} = a pure Nash equilibrium. Then:

$$\begin{aligned} \text{welfare}(\mathbf{b}) &= \text{Rev}(\mathbf{b}) + \sum_i u_i(\mathbf{b}) && [\text{defn of utility}] \\ &\geq \text{Rev}(\mathbf{b}) + \sum_i u_i(b_i^*, \mathbf{b}_{-i}) && [\mathbf{b} \text{ a Nash eq}] \\ &\geq \text{Rev}(\mathbf{b}) + [\lambda \cdot [\text{OPT Welfare}] - \text{Rev}(\mathbf{b})] \\ &= \lambda \cdot [\text{OPT Welfare}] \end{aligned}$$

Guarantee for FPAs

Step 1: fix valuation profile \mathbf{v} . (Full-information case.)

Technical claim: for suitable choice of \mathbf{b}^* , for every \mathbf{b} ,

$$\sum_i u_i(b_i^*, b_{-i}) \geq \frac{1}{2} \cdot [\text{OPT Welfare}] - \text{Revenue}(\mathbf{b}).$$

Proof: Set $b_i^* = v_i/2$ for every i . (a la [Lucier/Paes Leme 11])

- since LHS ≥ 0 , can assume $\frac{1}{2} \cdot [\max_i v_i] > \max_i b_i$
- suppose bidder 1 has highest valuation. Then:

$$u_1(b_1^*, b_{-1}) = v_1 - (v_1/2) = v_1/2 \geq \frac{1}{2} \cdot [\text{OPT Welfare}]$$

The Return of 63%

Step 1: fix valuation profile \mathbf{v} . (Full-information case.)

Technical claim: for suitable choice of \mathbf{b}^* , for every \mathbf{b} ,

$$\sum_i E[u_i(b_i^*, b_{-i})] \geq .63 \cdot [\text{OPT Welfare}] - \text{Revenue}(\mathbf{b}).$$

Proof: Draw b_i^* from $[0, .63v_i]$, density $1/(v_i - b)$ [Syrkanis 12]

- suppose bidder 1 has highest valuation. Then:

$$E[u_1(b_1^*, b_{-1})] \geq .63 \cdot [\text{OPT Welfare}] - \text{Revenue}(\mathbf{b})$$

Note: “private” baseline b_i^* depends on v_i but *not* on v_{-i} .

Guarantees for BNE

Assume: for suitable choice of $b_1^*(v), \dots, b_n^*(v)$, for every v and \mathbf{b} : $\sum_i u_i(b_i^*(v_i), \mathbf{b}_{-i}) \geq \lambda \cdot [\text{OPT Welfare}(v)] - \text{Rev}(\mathbf{b})$.

Claim: (\approx [Lucier/Paes Leme 11]) for all (possibly correlated) valuation distributions, POA of Bayes-Nash eq is $\geq \lambda$.

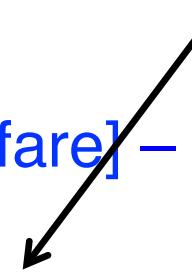
Proof: Let $\mathbf{b}()$ = a Bayes-Nash equilibrium. Then:

$$\begin{aligned} E_v[\text{welfare}(\mathbf{b}(v))] &= E_v[\text{Rev}(\mathbf{b}(v))] + \sum_i E_v[u_i(\mathbf{b}(v))] \quad [\text{defn of utility}] \\ &\geq E_v[\text{Rev}(\mathbf{b}(v))] + \sum_i E_v[u_i(b_i^*(v_i), \mathbf{b}_{-i}(v_{-i}))] \quad [\mathbf{b} \text{ a BNE}] \\ &\geq E_v[\text{Rev}(\mathbf{b}(v))] + [\lambda \cdot E_v[\text{OPT Welfare}(v)] - E_v[\text{Rev}(\mathbf{b}(v))]] \\ &= \lambda \cdot E_v[\text{OPT Welfare}(v)] \end{aligned}$$

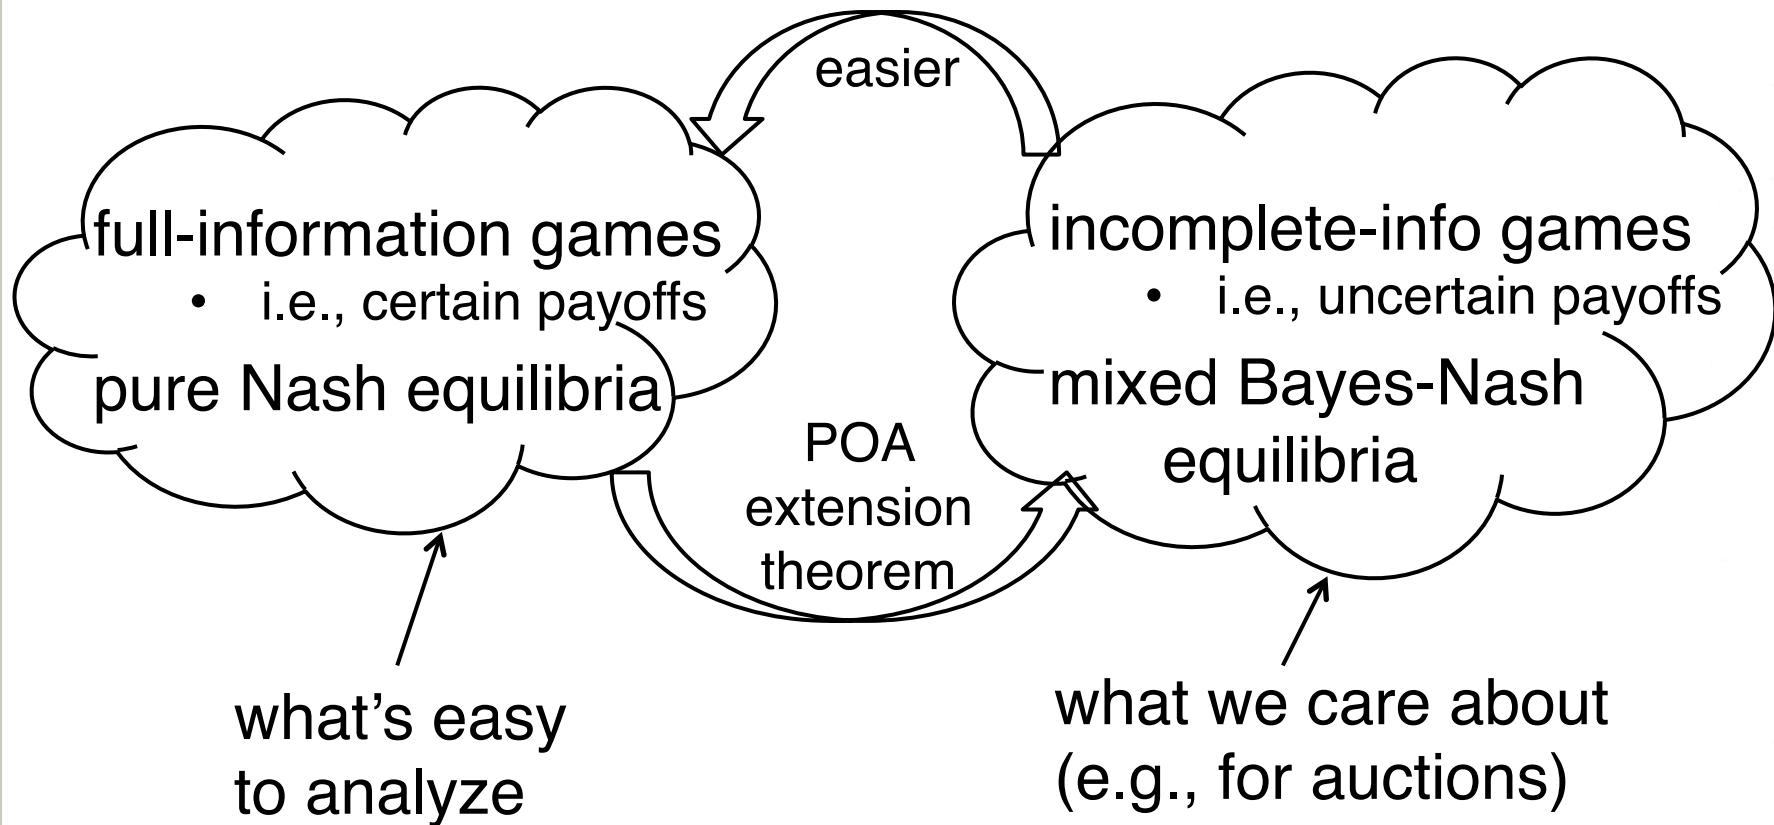
Smoothness Paradigm

1. Fix a setting and the private valuations \mathbf{v} .
 2. Choose private baseline strategies \mathbf{b}^* .
(i.e., b_i^* can depend on v_i but not on \mathbf{v}_{-i})
 3. Fix outcome \mathbf{b} .
 4. Prove $\sum_i u_i(b_i^*, \mathbf{b}_{-i}) \geq \lambda \cdot [\text{OPT Welfare}] - \text{Revenue}(\mathbf{b})$.
 5. Conclude that POA of Bayes-Nash equilibria is $\geq \lambda$.
- λ -smooth auction
- 

Smoothness Paradigm

1. Fix a setting and the private valuations \mathbf{v} .
 2. Choose private baseline strategies \mathbf{b}^* .
(i.e., b_i^* can depend on v_i but not on \mathbf{v}_{-i})
 3. Fix outcome \mathbf{b} .
 4. Prove $\sum_i u_i(b_i^*, \mathbf{b}_{-i}) \geq \lambda \cdot [\text{OPT Welfare}] - \text{Revenue}(\mathbf{b})$.
 5. Conclude that POA of Bayes-Nash equilibria is $\geq \lambda$.
- extension theorem*
- 

Extension Theorem (Informal)



[Roughgarden 09], [Lucier/Paes Leme 11], [Roughgarden 12],
[Syrkanis 12], [Syrkanis/Tardos 13], ...

The Story So Far

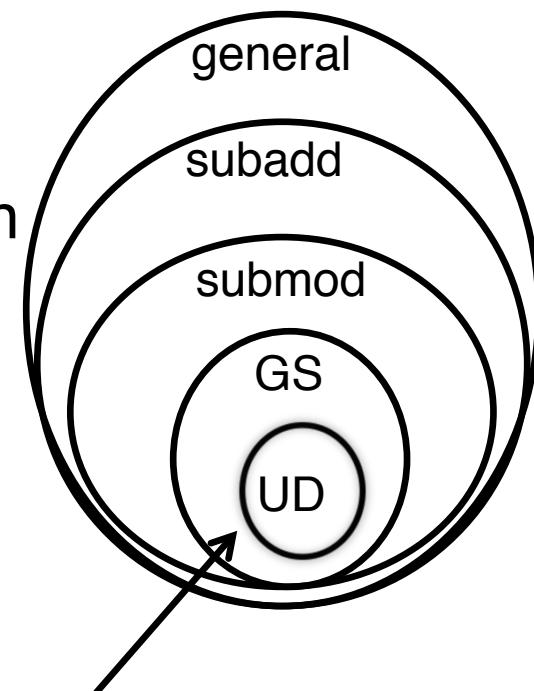
Summary: for all (possibly correlated) valuation distributions, every Bayes-Nash equilibrium of a first-price auction has welfare at least 63% of the maximum possible.

- 63% is tight for correlated valuations [Syrkanis 14]
- independent valuations = worst-case POA unknown
 - worst known example = 87% [Hartline/Hoy/Taggart 14]
- **next:** to what extent does 63% extend to simultaneous single-item auctions?
 - would be tight even with full-info [Engelbrecht-Wiggins/Weber 79]

Simultaneous First-Price Auctions (A Composition Theorem)

Multi-Item Auctions

- suppose m different items
- for now: *unit-demand* valuations
- each bidder i has private valuation v_{ij} for each item j
- $v_i(S) := \max_{j \text{ in } S} v_{ij}$
- *submodular*: if $j \in S \subseteq T$,
 $v_i(S + j) - v_i(S) \geq v_i(T + j) - v_i(T)$



you are here

Simultaneous Composition

- suppose have mechanisms M_1, \dots, M_m
- in their *simultaneous composition*:
 - new action space = product of the m action spaces
 - new allocation rule = union of the m allocation rules
 - new payment rule = sum of the m payment rules
- example: each M_j a single-item first-price auction

Question: as a unit-demand bidder, how should you bid?
(not so easy)

Does Composition Preserve Smoothness?

Hypothesis: every single-item auction M_j is λ -smooth: there exist private baselines $b^*_1(), \dots, b^*_n()$ such that, for every v and b :

$$\sum_i u_i(b^*_i, b_{-i}) \geq \lambda \cdot [\text{OPT Welfare}(v)] - \text{Rev}(b).$$

Question: is composed mechanism also λ -smooth?

Proof idea: Fix unit-demand valuations v , fixes OPT.

- baseline strategy for a bidder i that gets item j in OPT
 - bid 0 in mechanisms other M_j
 - in M_j , use assumed baseline strategy for M_j

Public vs. Private Baselines

First-price auction: set $b_i^* = v_i/2$ for every i .

- independent of v_{-i} (*private* baseline strategies)

Simultaneous first-price auctions: proposed b_i^* is “bid half your value only on the item j you get in $\text{OPT}(v)$.”

- *public* baseline strategies
- not well defined unless v_{-i} known

Extension Theorem (BNE)

Assume: there exist private baselines $b_1^*(\cdot), \dots, b_n^*(\cdot)$ such that, for every v and b :

$$\sum_i u_i(b_i^*, b_{-i}) \geq \lambda \cdot [\text{OPT Welfare}(v)] - \text{Rev}(b).$$

Claim: (\approx [Lucier/Paes Leme 11]) for all (possibly correlated) valuation distributions, POA of Bayes-Nash eq is $\geq \lambda$.

Proof: Let $b(\cdot)$ = a Bayes-Nash equilibrium. Then:

$$\begin{aligned} E_v[\text{welfare}(b(v))] &= E_v[\text{Rev}(b(v))] + \sum_i E_v[u_i(b(v))] \quad [\text{defn of utility}] \\ &\geq E_v[\text{Rev}(b(v))] + \sum_i E_v[u_i(b_i^*(v_i), b_{-i}(v_{-i}))] \quad [b \text{ a BNE}] \\ &\geq E_v[\text{Rev}(b(v))] + [\lambda \cdot E_v[\text{OPT Welfare}] - E_v[\text{Rev}(b(v))]] \\ &= \lambda \cdot E_v[\text{OPT Welfare}] \end{aligned}$$

deviation can depend
on v_i but not v_{-i}

Counterexample

Fact: [Feldman/Fu/Gravin/Lucier 13], following [Bhawalkar/Roughgarden 11] there are (highly correlated) valuation distributions over unit-demand valuations such that every Bayes-Nash equilibrium has expected welfare arbitrarily smaller than the maximum possible.

- idea: plant a random matching plus some additional highly demanded items; by symmetry, a bidder can't detect the item "reserved" for it

Does Composition Preserve Smoothness?

Hypothesis: every single-item auction M_j is λ -smooth
with public baselines: for every v , there exists baselines b^* such that, for every b :

$$\sum_i u_i(b_i^*, b_{-i}) \geq \lambda \cdot [\text{OPT Welfare}(v)] - \text{Rev}(b).$$

Theorem: [Syrkanis/Tardos 13] with unit-demand bidders, composed mechanism is λ -smooth w/public baselines.

- holds for arbitrary smooth M_j 's, submodular valuations

Proof idea: Baseline strategy for a bidder i that gets item j in OPT: use baseline strategy for M_j , 0 on other items.

Modified Extension Theorem

Hypothesis: mechanism is λ -smooth with public baselines: for every \mathbf{v} , there exists baselines \mathbf{b}^* such that, for every \mathbf{b} :

$$\sum_i u_i(b_i^*, \mathbf{b}_{-i}) \geq \lambda \cdot [\text{OPT Welfare}(\mathbf{v})] - \text{Rev}(\mathbf{b}).$$

Theorem: [Syrkanis/Tardos 13], following [Christodoulou/Kovacs/Schapira 08] for all product valuation distributions, POA of Bayes-Nash eq is $\geq \lambda$.

Proof idea: transform public baseline to interim deviation:

- sample \mathbf{w}_{-i} from prior distribution
- play baseline strategy for valuation profile $(\mathbf{v}_i, \mathbf{w}_{-i})$

A Promise Fulfilled

Consequence: for all *product* unit-demand (or submodular) valuation distributions, every Bayes-Nash equilibrium of simultaneous first-price auctions has welfare at 63% of the maximum possible.

Proof sketch:

1. a single-item first-price auction is .63-smooth
2. composition => simultaneous FPAs also .63-smooth
with public baseline strategies
3. *modified* extension theorem => Bayes-Nash POA of simultaneous FPAs $\geq .63$ (any *product* distribution)

When Are Equilibria of Simple Auctions Near-Optimal?

Part 2: Impossibility Results

Tim Roughgarden (Stanford)

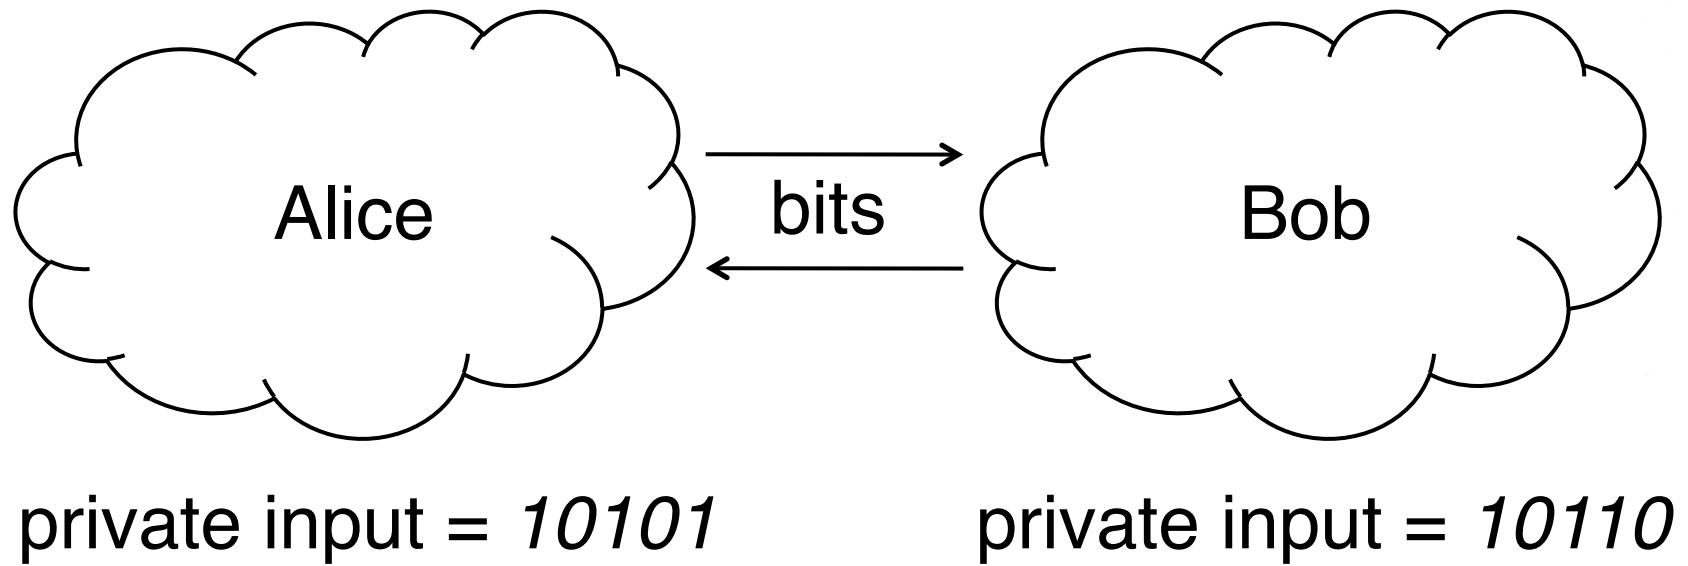
*Reference: Roughgarden/Syrgkanis/Tardos,
“The Price of Anarchy in Auctions,” Journal of
Artificial Intelligence Research (JAIR)*

Communication Complexity

General references:

- Kushilevitz/Nisan, *Communication Complexity*, 1996
- Roughgarden, *Communication Complexity (for Algorithm Designers)*, 2016.

Are Two Strings Identical?

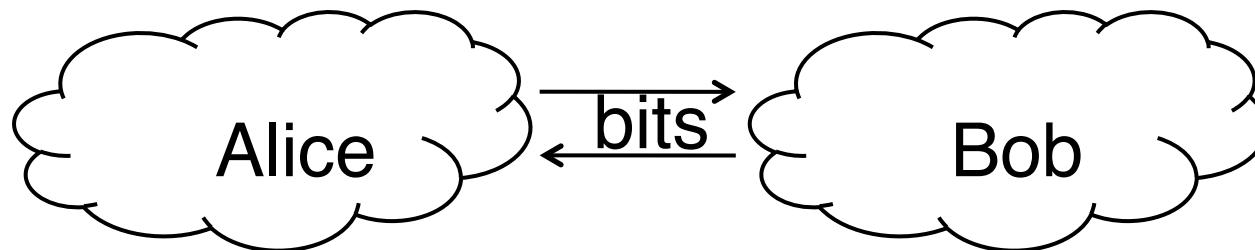


Communication Complexity (Deterministic)

- fix $f:\{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ (e.g., EQUALITY)
- Alice, Bob agree on “communication protocol”
 - specifies who sends what bit (0 or 1) when
- each sent bit depends *only* on that player’s private input and the history-so-far of protocol (for now, deterministic)
- at end of protocol, at least one player should know the result $f(x,y)$ [x = Alice’s input, y = Bob’s input]
- cost of protocol = max # of bits sent (for worst-case x,y)
- *communication complexity of f* = min-cost of any protocol

Example: EQUALITY

- deterministic communication complexity (DCC) of EQUALITY is $\leq n$ ($n = \text{length of } x \text{ and } y$)
 - Alice can always send her entire input to Bob
- intuition: $\text{DCC}(\text{EQUALITY}) \geq n$ as well



private input = 10101

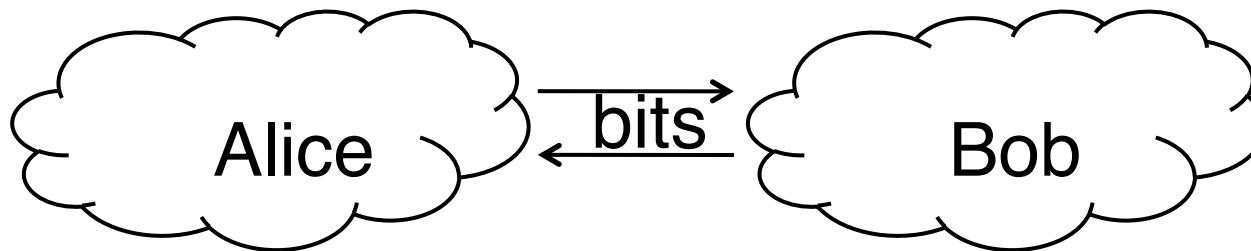
private input = 10100

- issue: communication protocols can be very clever!

The Cost of Convincing

Question: if I know x and y :

- how hard is it to convince Alice & Bob that $x \neq y$?
(if indeed $x \neq y$)
- what about that $x = y$? (if indeed $x = y$)



private input = *10101*

private input = *10100*

Communication Complexity (Nondeterministic)

- Alice, Bob want to evaluate f on private inputs x,y
- an omniscient third party (the “prover” P) knows x,y
- nondeterministic communication protocol:
 - prover writes an alleged proof that $f(x,y)=b$ in public view
 - Alice, Bob simultaneously decide whether to accept or reject the proof (as a function only of private input and the prover’s proof)
- correctness: never the case that Alice, Bob accept proof that $f(x,y)=b$ when in fact $f(x,y)\neq b$ [also, “completeness”]
- cost of protocol = max proof length, in bits (for worst x,y)
- *nondeterministic CC of f* = min-cost of any protocol

EQUALITY Resolved

Theorem: [Yao 79] The nondeterministic communication complexity of EQUALITY is n (to prove that $f(x,y)=1$).

- automatically implies same for deterministic case

Lemma: Let π be an alleged proof. Suppose Alice, Bob both accept when inputs are (x_1, y_1) or (x_2, y_2) . Then Alice, Bob also accept π when inputs are (x_1, y_2) or (x_2, y_1) .

- accept/reject decisions depend only on private input and proof π

EQUALITY Resolved

Lemma: Let π be an alleged proof. Suppose Alice, Bob both accept when inputs are (x_1, y_1) or (x_2, y_2) . Then Alice, Bob also accept π when inputs are (x_1, y_2) or (x_2, y_1) .

Corollary: For every $w \neq z$, there is no proof π such that Alice, Bob accept π for inputs (w, w) and (z, z) .

- if there were, π would also convince Alice, Bob that $f(w, z)$ and $f(z, w) = 1$ (contradicts correctness)

Corollary: Prover must use different proof for each of the 2^n inputs of form (w, w) . [“fooling set”]

- \Rightarrow at least one proof uses $\geq n$ bits [proves Theorem]

Communication Complexity of Welfare-Maximization

Distinguishing High vs. Low Social Welfare

Welfare-maximization problem: n players.

- **private inputs:** valuations v_i over set of m items
- **goal:** evaluate function f where: [for known W, α]
 - $f(v_1, \dots, v_n) = 1$ if there is an allocation with welfare $\geq W$
 - $f(v_1, \dots, v_n) = 0$ if there is no allocation with welfare $\leq \alpha W$
 - $f(v_1, \dots, v_n)$ undefined otherwise
- protocol only need to be correct on inputs where f is defined (can behave arbitrarily otherwise)
 - only easier than computing the maximum-welfare allocation
 - impossibility results only stronger

Nondeterministic Protocols

- nondeterministic communication protocol: prover writes an alleged proof π that $f(v_1, \dots, v_n) = b$ in public view (b in $\{0,1\}$)
- all players simultaneously decide to accept or reject the proof (as a function only of own input and π)
- correctness: never the case that all players accept proof that $f(v_1, \dots, v_n) = b$ for an input with $f(v_1, \dots, v_n) = 1 - b$
 - [also, “completeness”]
- **note:** when $f(v_1, \dots, v_n) = 1$, easy to prove it
 - prover writes down allocation S_1, \dots, S_n and players’ values $v_1(S_1), \dots, v_n(S_n)$, players can check legitimacy
 - proof length = scales roughly linearly with n, m

Known Lower Bounds

Theorem: [Nisan 02] For every constant $a > 0$, for all large enough m , every nondeterministic protocol differentiating

$$\text{OPT welfare}(\mathbf{v}) \geq W \text{ vs. } \text{OPT welfare}(\mathbf{v}) \leq aW$$

for general valuations \mathbf{v} has cost (i.e., max proof length) exponential in m .

Theorem: [Dobzinski/Nisan/Schapira 05] For every constant $a > \frac{1}{2}$, for all large enough m , every nondeterministic protocol differentiating

$$\text{OPT welfare}(\mathbf{v}) \geq W \text{ vs. } \text{OPT welfare}(\mathbf{v}) \leq aW$$

for subadditive valuations \mathbf{v} has cost exponential in m .

Optimal Simple Auctions

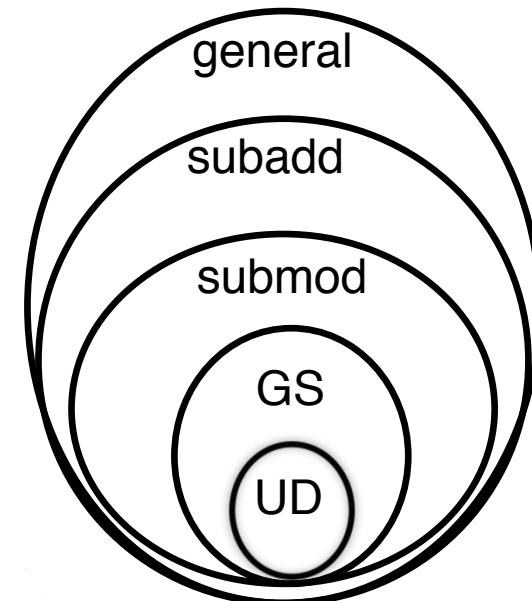
Tight POA Bounds

Theorem: [Feldman/Fu/Gravin/Lucier 13], [Christodoulou/Kovacs/Sgouritsa/Tang 14] the worst-case POA of simultaneous first-price auctions with subadditive bidder valuations is precisely 50%.

- Bayes-Nash equilibria,
arbitrary product prior

subadditive valuations:

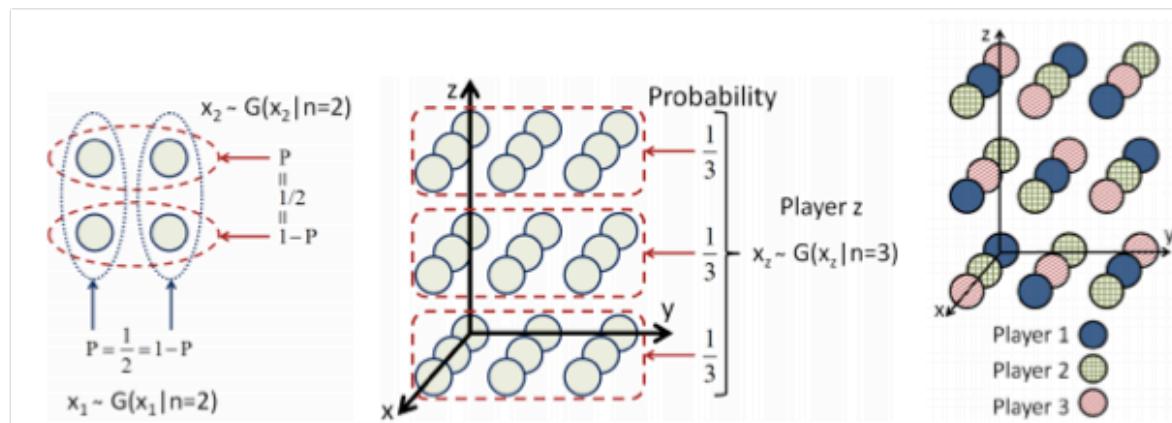
- $v_i(A \cup B) \leq v_i(A) + v_i(B)$ for
all disjoint A, B



Tight POA Bounds

Theorem: [Feldman/Fu/Gravin/Lucier 13], [Christodoulou/Kovacs/Sgouritsa/Tang 14] the worst-case POA of simultaneous first-price auctions with subadditive bidder valuations is precisely 50%.

Explicit
Lower
Bound:



Tight POA Bounds

Theorem: [Feldman/Fu/Gravin/Lucier 13], [Christodoulou/Kovacs/Sgouritsa/Tang 14] the worst-case POA of simultaneous first-price auctions with subadditive bidder valuations is precisely 50%.

Question: Can we do better?

(without resorting to the VCG mechanism)

The Upshot

Meta-theorem: equilibria are generally bound by the same limitations as algorithms with polynomial computation or communication.

- lower bounds without explicit constructions!

Caveats: requires that equilibria are

- guaranteed to exist (e.g., mixed Nash equilibria)
- can be efficiently verified

Example consequence: no “simple” auction has POA $> 50\%$ for bidders with subadditive valuations.

From Protocol Lower Bounds to POA Lower Bounds

Main Theorem: [Roughgarden 14] Suppose:

- no nondeterministic subexponential-communication protocol approximates the welfare-maximization problem (with valuations V) to within factor of a .
 - i.e., impossible to decide $\text{OPT} \geq W$ vs. $\text{OPT} \leq aW$

Then worst-case POA of ε -approximate mixed Nash equilibria of every “simple” mechanism is at most a .

- simple = number of strategies sub-doubly-exponential in m
- ε can be as small as inverse polynomial in n and m

Point: : reduces impossibility results for equilibria to impossibility results for communication protocols.

Consequences (I)

Theorem: [Dobzinski/Nisan/Schapira 05] For every constant $\alpha > \frac{1}{2}$, for all large enough m , every nondeterministic protocol differentiating

$$\text{OPT welfare}(\mathbf{v}) \geq W \text{ vs. } \text{OPT welfare}(\mathbf{v}) \leq \alpha W$$

for subadditive valuations \mathbf{v} has cost exponential in m .

Corollary: (via Main Theorem) simultaneous first-price auctions is an *optimal* simple mechanism for subadditive valuations.

- best-possible worst-case POA (of ϵ -equilibria)

Consequences (II)

Theorem: [Nisan 02] For every constant $\alpha > 0$, for all large enough m , every nondeterministic protocol differentiating

$$\text{OPT welfare}(\mathbf{v}) \geq W \text{ vs. } \text{OPT welfare}(\mathbf{v}) \leq \alpha W$$

for general valuations \mathbf{v} has cost (i.e., max proof length) exponential in m .

Corollary: (via Main Theorem) *no* simple mechanism has a non-trivial (i.e., constant-factor) welfare guarantee for bidders with general valuations.

- complexity (e.g., package bidding) appears necessary

Practical Rules of Thumb

Folklore belief #1: without strong complements, simple auctions work pretty well.

- loss in outcome quality appears small
- demand reduction exists, but not a dealbreaker

Folklore belief #2: with strong complements, simple auctions aren't good enough.

- *loss in outcome quality could be big*
- *exposure problem exists, and is a dealbreaker*

Proof of Main Theorem

Why Approximate MNE?

Issue: in simultaneous first-price auctions, number of strategies per-player $N = (V_{\max} + 1)^m$ [exponential in m]

- valuations, bids assumed integral and poly-bounded
- can't efficiently communicate a MNE.

Theorem: [Lipton/Markakis/Mehta 03] a game with n players and N strategies per player has an ε -approximate mixed Nash equilibrium with support size polynomial in n , $\log N$, and ε^{-1} .

- proof idea based on sampling from an exact MNE

Proof of Theorem

Suppose worst-case POA of ε -MNE is $\rho > \alpha$:

Input: v_i 's s.t.
either (i) $OPT \geq W$ or (ii)
 $OPT \leq \alpha W$



Protocol:
“proof” =
 ε -MNE x with
small support
(exists by
LMM); players
verify it privately

if $E[wel(x)] > \alpha W$
then $OPT > \alpha W$
so in case (i)

if $E[wel(x)] \leq \alpha W$
then $OPT \leq (\alpha / \rho)W < W$
so in case (ii)

Key point: every ε -MNE is a short, privately
verifiable certificate for membership in case (ii).

Two Open Questions

1. Tight POA bounds for important auction formats
 1. e.g. first-price auctions, independent valuations (63% vs. 87%)
2. Best simple auction for submodular valuations?
 1. simultaneous first-price auctions give 63% [Syrgkanis/Tardos 13], [Christodoulou et al 14]
 2. > 77% impossible with a simple auction [Dobzinski/Vondrak 13] + [Roughgarden 14]
 3. > 63% *is* possible with only polynomial communication [Feige/Vondrak 06]