Mechanism design via extensive-form games

Brian Zhang

Problem 1: Stackelberg equilibria in normal-form games

		Follower X	Follower Y
Leader	X	1, 1	11, 0
	Y	0, 0	10, 1

unique Nash (1, 1)

Idea: If P1 can **commit** to playing X w.p. 1/2-ε, then:

- P2's BR is to play Y
- P1 gets value ≈10.5

Problem 1: Stackelberg equilibria in normal-form games

		Follower	
		X	Y
Leader	X	1, 1	11, 0
	Y	0, 0	10, 1

Idea: If P1 can **commit** to playing X w.p. 1/2, then:

- P2's BR is to play Y
- P1 gets value 10.5

We'll ignore tiebreaking.

Equivalently:

- P1 issues a **recommendation** (here Y) to P2
- P2 must satisfy an **obedience constraint**

Also equivalently: Optimal equilibrium for P1, ignoring P1's incentive constraint

$$\max_{\substack{x_L \in \Delta(A_L) \\ x_F \in \Delta(A_F)}} u_L(x_L, x_F)$$

$$\text{s.t.} \quad u_F(x_L, x_F) \geq u_F(x_L, a_F) \quad \forall a_F \in A_F$$

~~$u_L(x_L, x_F) \geq u_L(a_L, x_F) \quad \forall a_L \in A_L$~~

(with this constraint, it would be optimal Nash eq.)

Problem 1: Stackelberg equilibria in normal-form games

		Follower X	Follower Y
Leader	X	1, 1	11, 0
	Y	0, 0	10, 1

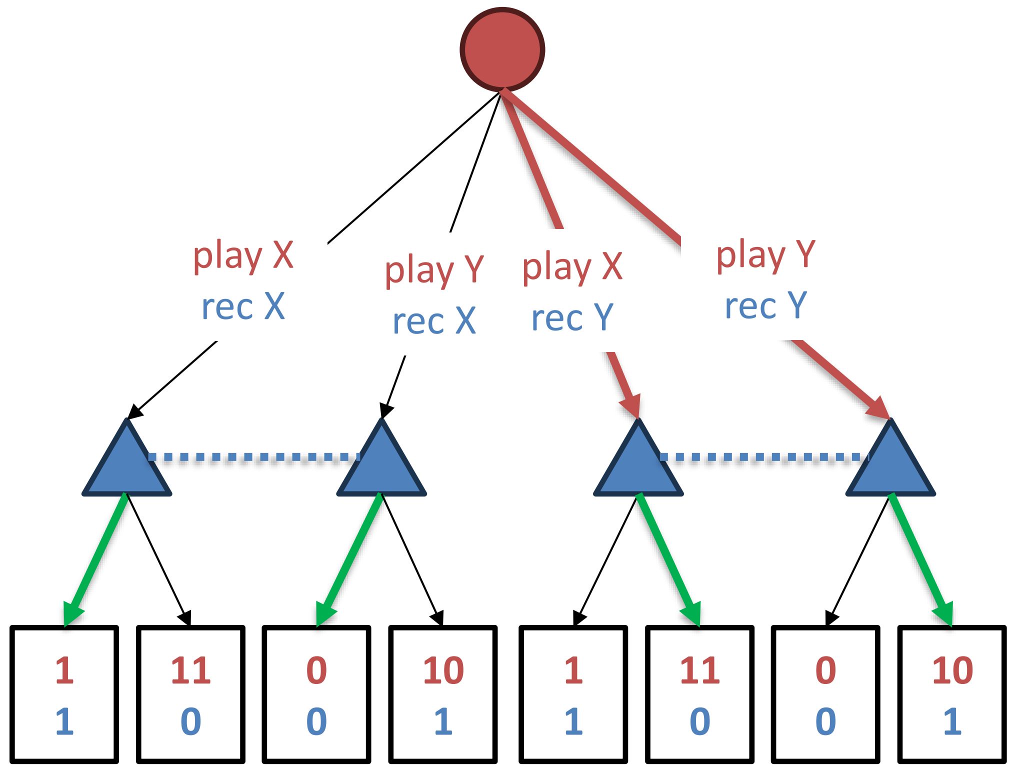

Stackelberg equilibrium = strategy for leader
s.t. holding leader's strategy fixed,
direct strategy is a **best response** for follower

Problem 2: Optimal correlated equilibria (for normal-form games)

Chicken

	Stop	Go
Stop	0, 0 0	0, 1 p
Go	1, 0 1-p	-5, -5 0

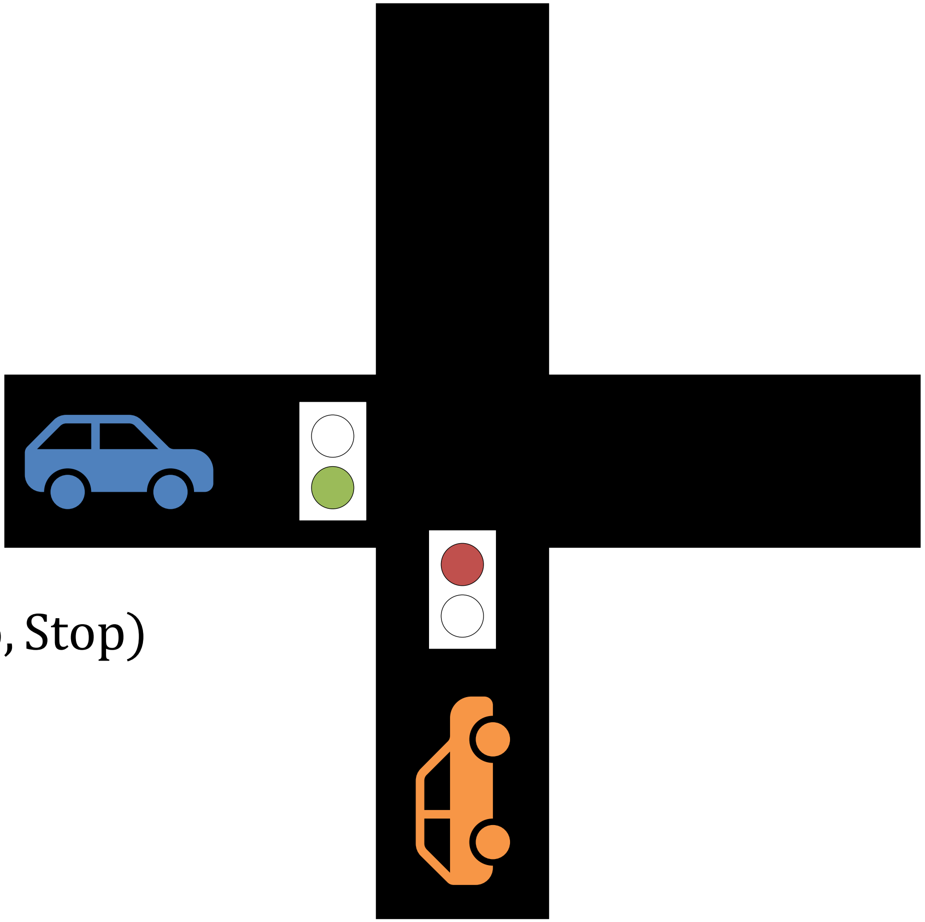

$$\mu = p \cdot (\text{Stop}, \text{Go}) + (1 - p) \cdot (\text{Go}, \text{Stop})$$

is a CE

Player strategies are **direct**: incentivized to follow recommendations

Problem 2: Optimal correlated equilibria (for normal-form games)

Chicken

	Stop	Go
Stop	0, 0 0	0, 1 p
Go	1, 0 1-p	-5, -5 0

Correlated equilibrium = **strategy for the mediator** s.t. holding mediator's strategy fixed, **direct profile is Nash equilibrium** for other players

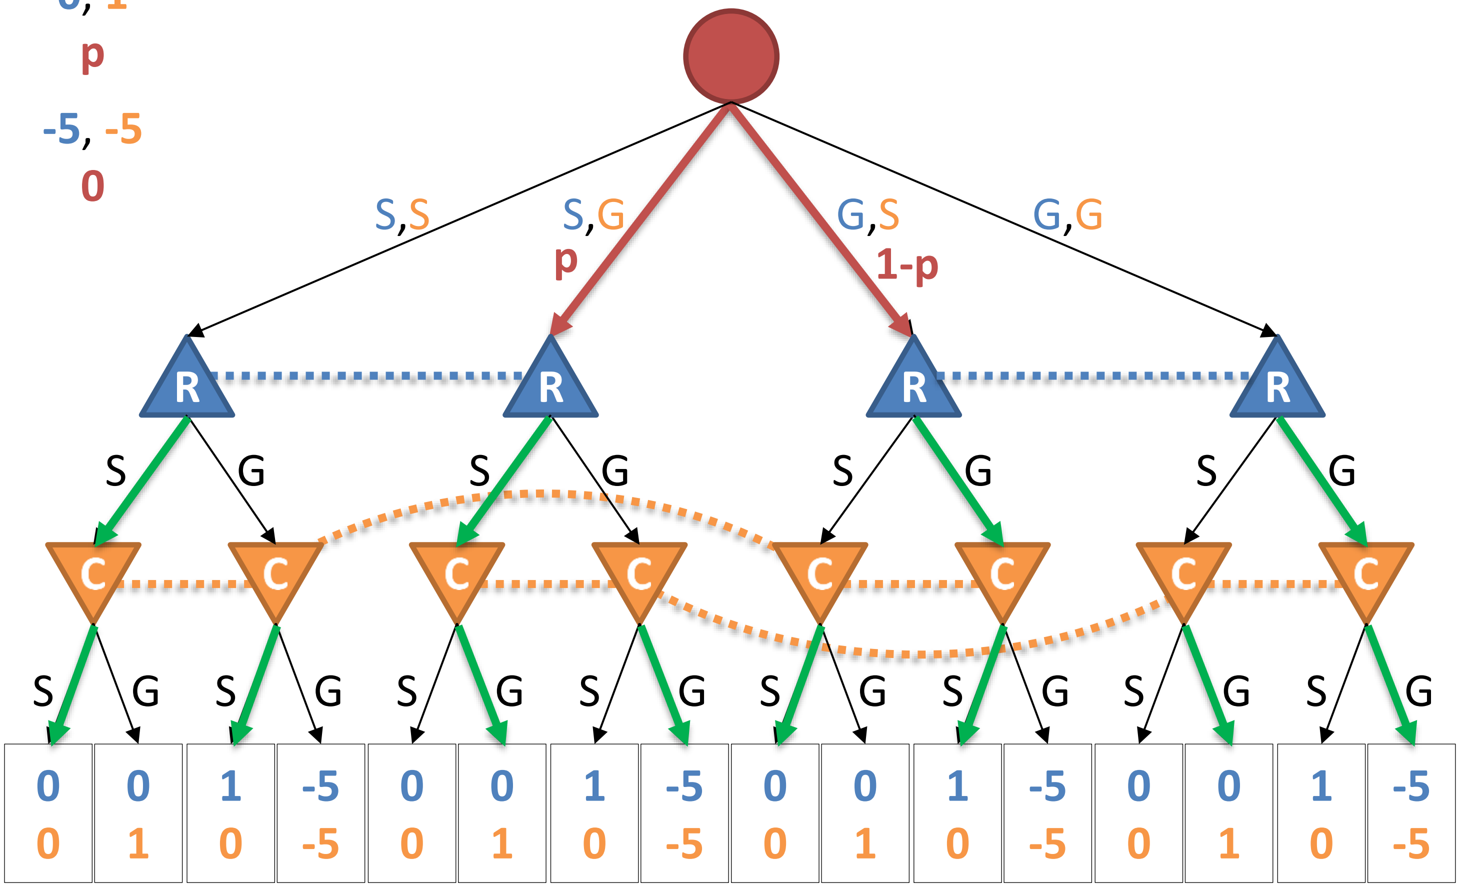

Problem 3: Mechanism design

Problem 3: Mechanism design

Theorem (Myerson, *Math of OR*'81):

Assuming buyer valuations are drawn i.i.d. from some distribution D, there exists **reserve price** r (dependent on D) for which the following mechanism is revenue-maximizing:

Auctioneer

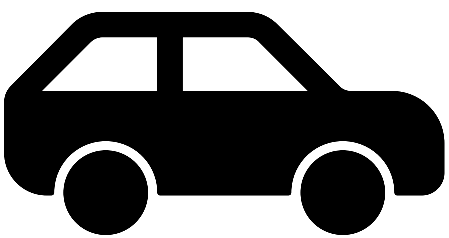

Buyer 1

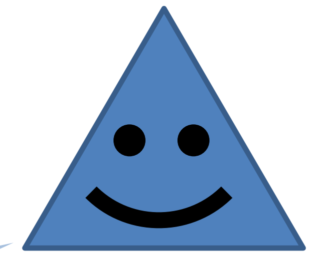

What are your valuations ("bids")?

$6k

$4k

$6k > r?

Yes

Highest bidder wins

Pays $\max($second-highest bid$, r)$

No

Nobody wins

Buyer 2

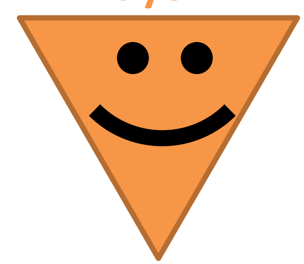

This mechanism is "direct": buyers bidding true values is Nash eqm

Problem 3: Mechanism design

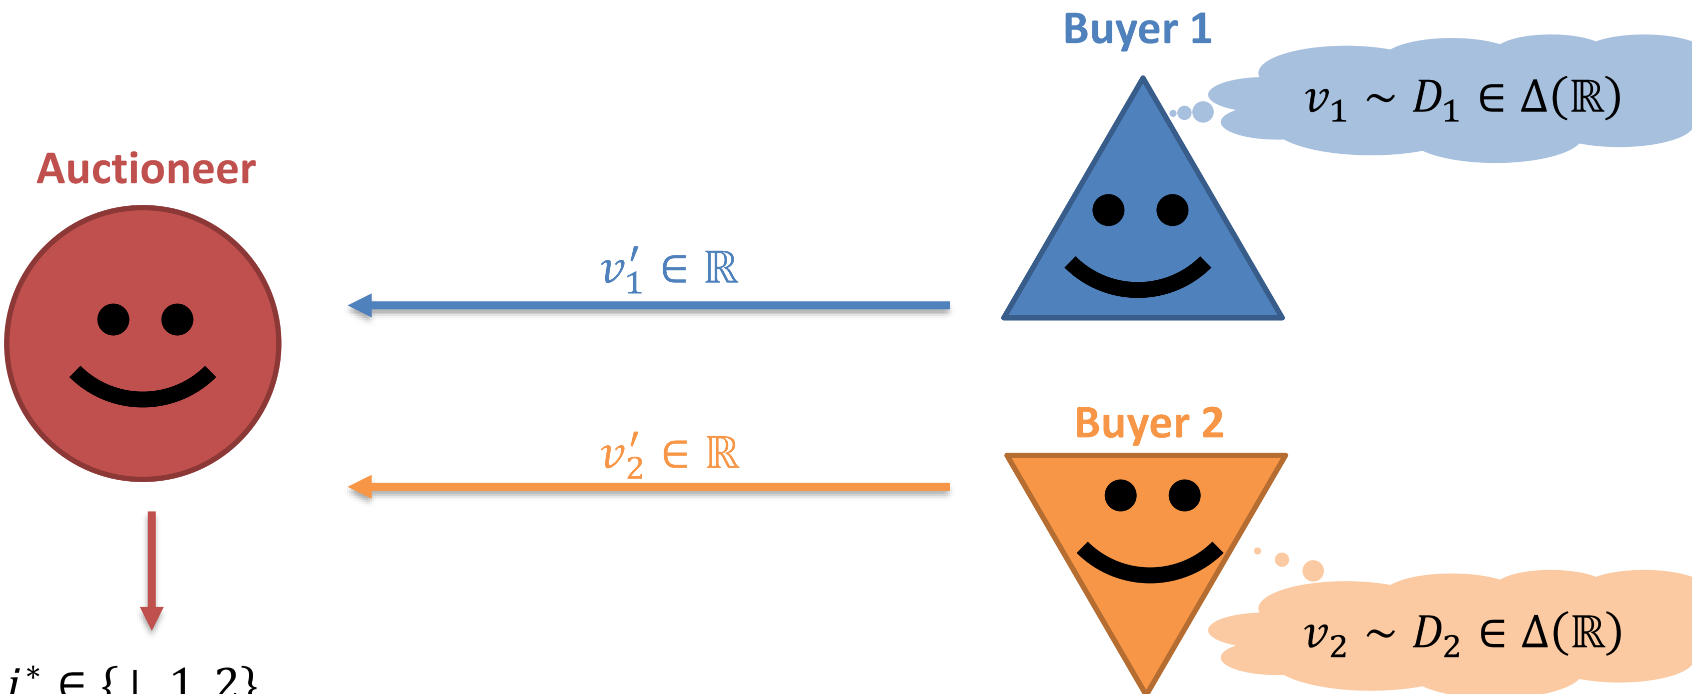

winner $i^* \in \{\perp, 1, 2\}$
price $p \in \mathbb{R}$

if $i^* = \perp$: everyone gets 0
else:

- $u_A = p, u_{i^*} = v_{i^*} - p$
- everyone else gets 0

Direct strategy of buyer i: report $v_i' = v_i$

Optimal mechanism = **strategy for the auctioneer**
s.t. holding auctioneer's strategy fixed,
direct profile is Nash equilibrium for other players

Problem 4: Information design

a.k.a. (Bayesian) persuasion

("Mechanism design, but backwards")

Kamenica & Gentzkow (*American Econometric Review*'11)

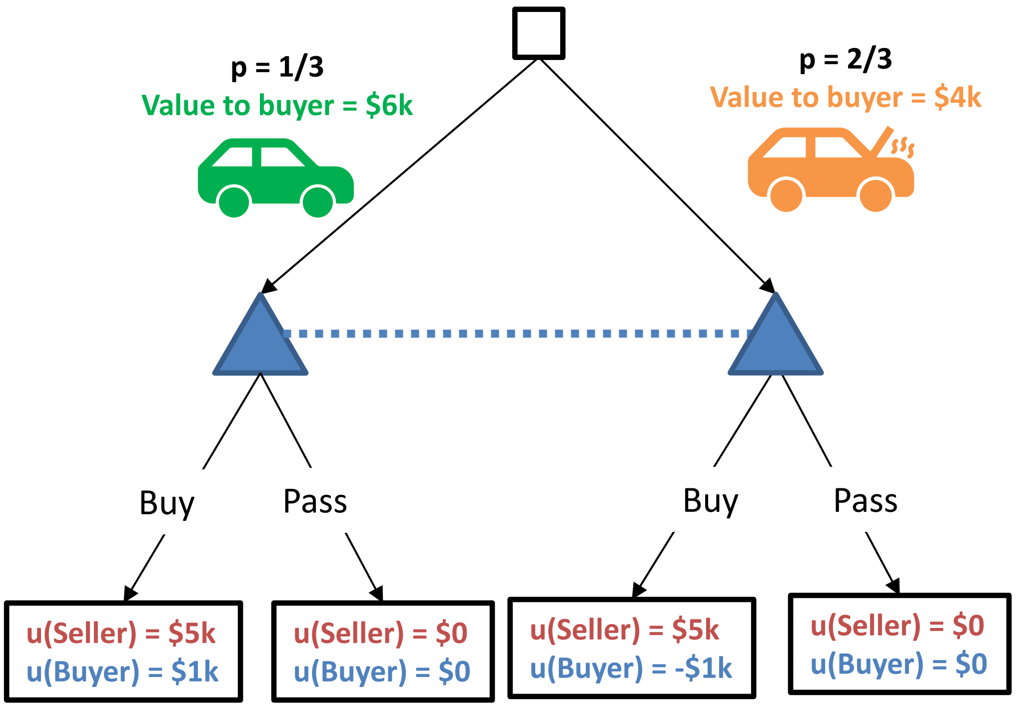

A car is on sale for $5,000…

…but only the seller knows the car's true quality

As before: The seller can commit, and send messages

Problem 4: Information design

a.k.a. (Bayesian) persuasion
("Mechanism design, but backwards")
Kamenica & Gentzkow (*American Econometric Review*'11)

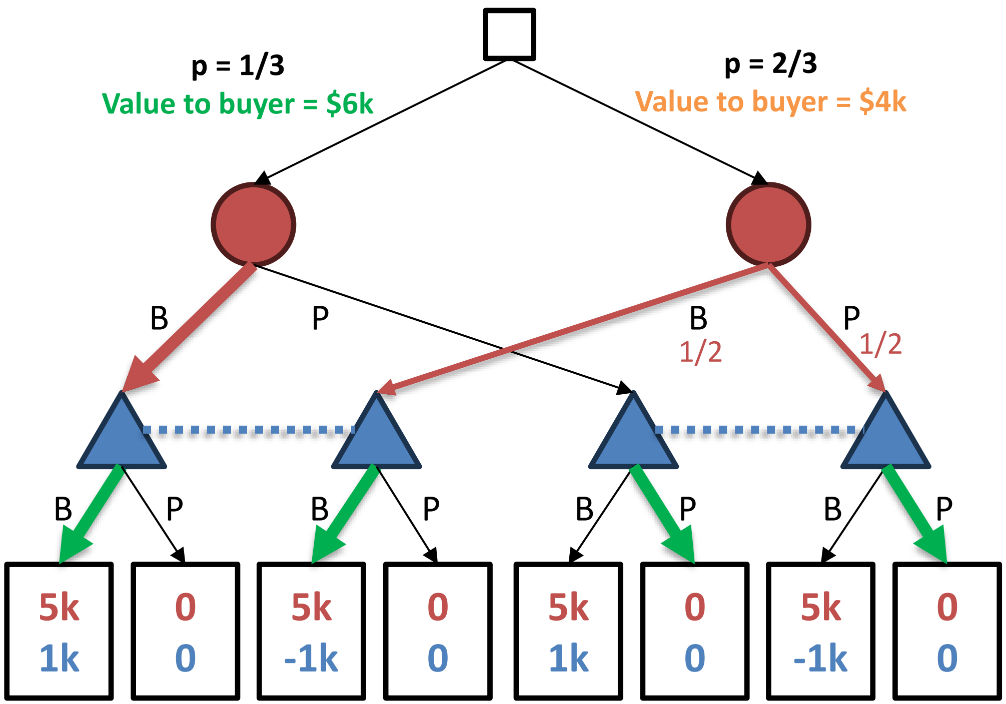

A car is on sale for $5,000…

…but only the seller knows the car's true quality

As before: The seller can commit, and send messages

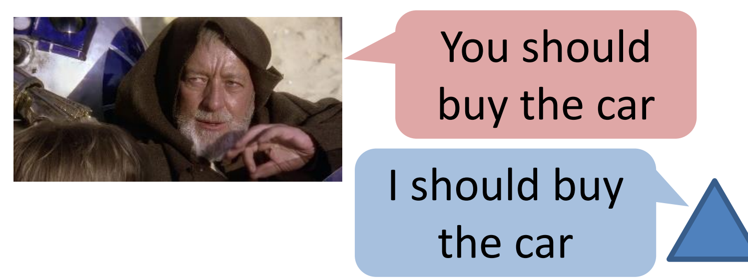

Buyer: Pr[car is good | seller says "buy"] = 1/2

⇒ Buyer's best response is to do what seller says (Strategy is **direct**)

Even though the car is good only 1/3 of the time, the seller sells the car 2/3 of the time!

Problem 4: Information design

a.k.a. (Bayesian) persuasion

("Mechanism design, but backwards")

Kamenica & Gentzkow (*American Econometric Review*'11)

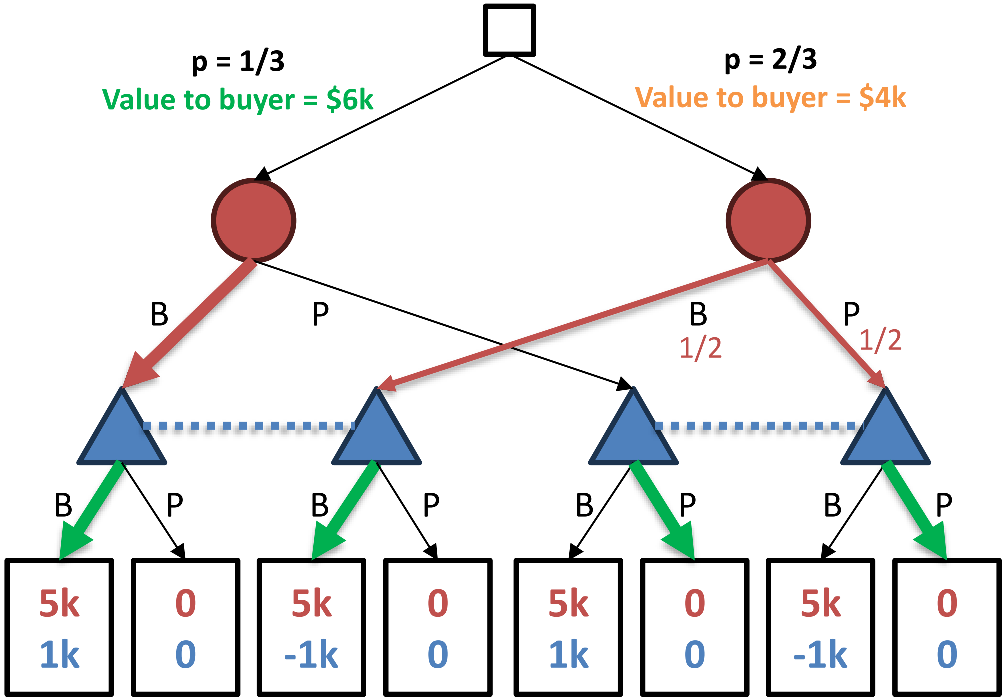

A car is on sale for $5,000…

…but only the seller knows the car's true quality

As before: The seller can commit, and send messages

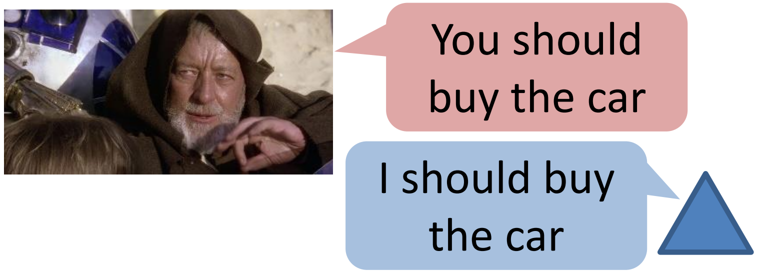

Optimal solution = **strategy for the seller**
s.t. holding seller's strategy fixed,
direct strategy is best response for the buyer

What's common to all these problems?

- **Optimization:** The mediator (leader/seller/correlation device) has some objective that it wants to optimize
- **Commitment:** The mediator commits to a strategy μ
- **Communication:** The mediator communicates with the players (gives them information/recommendations, or gets information from them).
  - Communication has no direct effect on the game; only purpose is to exchange information
  - Communication is **structured:** in all examples so far, it has been **information reports or action recommendations**

Rest of this lecture:

- How general is this?
- Can we compute these optimal mediator strategies efficiently?

Extensive-Form Games and Communication Equilibria

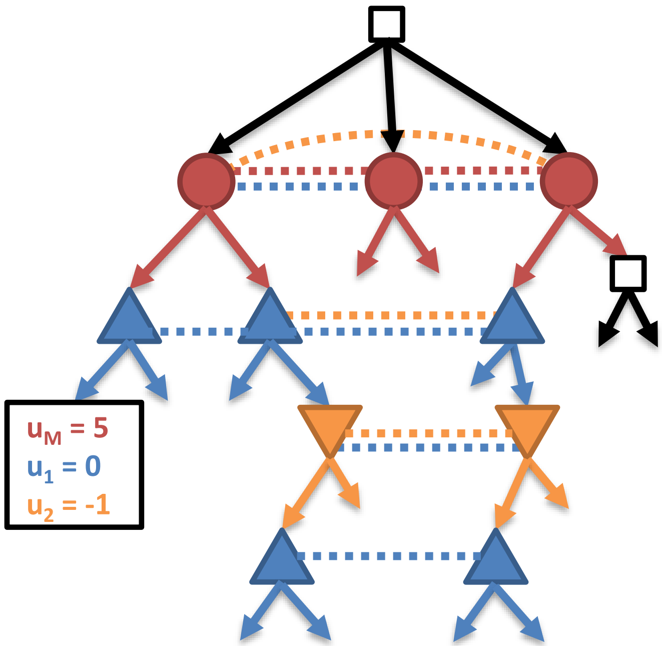

Communication is implicit. At every timestep in the game…

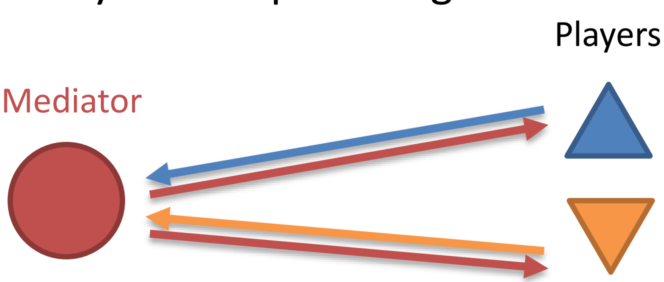

Definition:

Communication equilibrium: *tuple of (possibly randomized) strategies* (μ, x_1, x_2) s.t. all **players** (not incl. mediator) are best-responding:

$$u_1(\mu, x_1, x_2) = \max_{x_1'} u_1(\mu, x_1', x_2)$$

$$u_2(\mu, x_1, x_2) = \max_{x_2'} u_1(\mu, x_1, x_2')$$

Equivalently: (x_1, x_2) is a Nash equilibrium with μ held fixed

Extensive-Form Games and Communication Equilibria

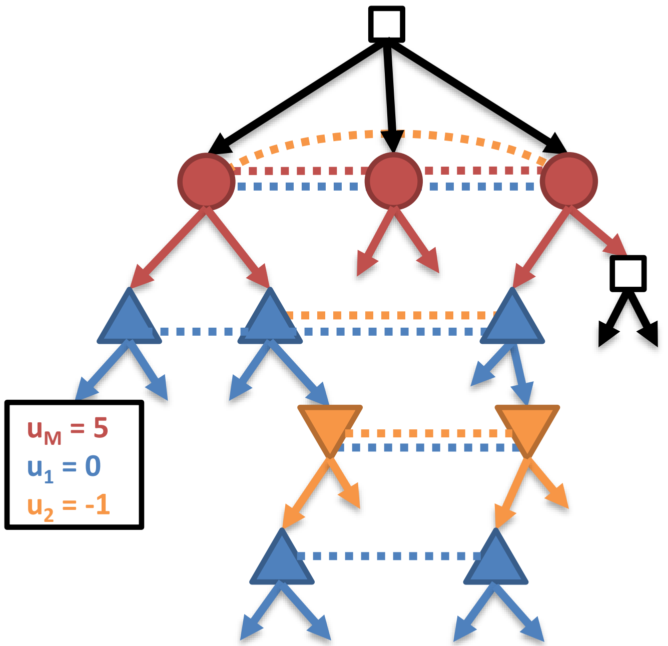

Definition:

Communication equilibrium: *tuple of (possibly randomized) strategies* (μ, x_1, x_2) s.t. all **players** (not incl. mediator) are best-responding:

$$u_1(\mu, x_1, x_2) = \max_{x_1'} u_1(\mu, x_1', x_2)$$

$$u_2(\mu, x_1, x_2) = \max_{x_2'} u_1(\mu, x_1, x_2')$$

Communication is implicit. At every timestep in the game…

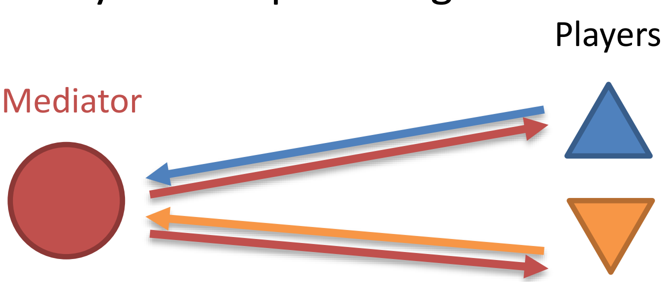

Main theorem (Zhang & Sandholm, NeurIPS'22):

There exists **poly(size of game tree)** algorithm that computes a communication equilibrium (μ, x_1, x_2) maximizing mediator's objective

$$u_M(\mu, x_1, x_2)$$

Extensive-Form Games and Communication Equilibria

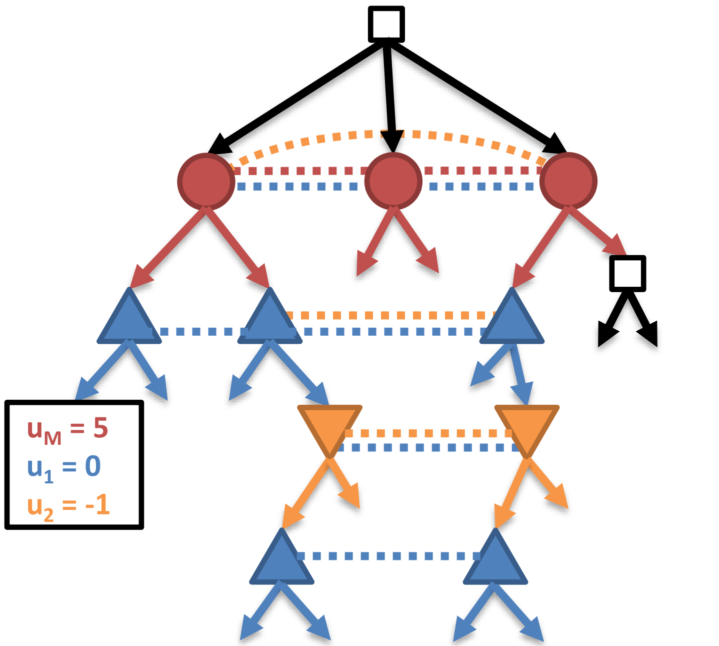

Definition:

Communication equilibrium: *tuple of (possibly randomized) strategies* (μ, x_1, x_2) s.t. all **players** (not incl. mediator) are best-responding:

$$u_1(\mu, x_1, x_2) = \max_{x_1'} u_1(\mu, x_1', x_2)$$

$$u_2(\mu, x_1, x_2) = max_{x_2'} u_1(\mu, x_1, x_2')$$

What are strategies?

Communication is implicit. At every timestep in the game…

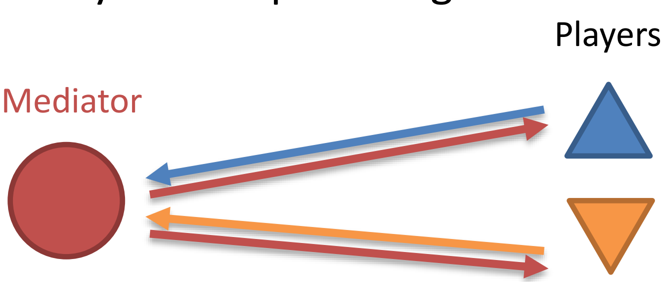

"If I observe t, then I should send m_1.

Then if I receive m_2, I should play action a, but if I receive m_3, I should play a'

Then, if…

Problem: Message space is infinite

Proof in Three Steps

- Step 1: Reduce game from infinite to finite.

  ("Revelation principle" introduces **structure** to the messages)

- Step 2: Reduce game from finite to polynomial.

  (Using more "without loss of generality" reductions)

- Step 3: Solve game.

  (LP duality)

Step 1: Revelation Principle

Theorem (Revelation Principle, *informal*)

(Forges, *Econometrica*'85, generalized in our full paper)

For every comm eqm, exists equivalent **direct** comm eqm. "Direct" means both:

1. Players' messages to mediator are **reports of private information**. In equilibrium, players always send their true private information
2. Mediator's messages to players are **action recommendations.** In equilibrium, players play actions that they are recommended

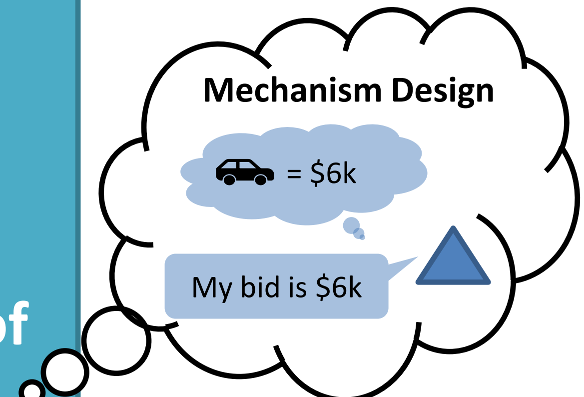

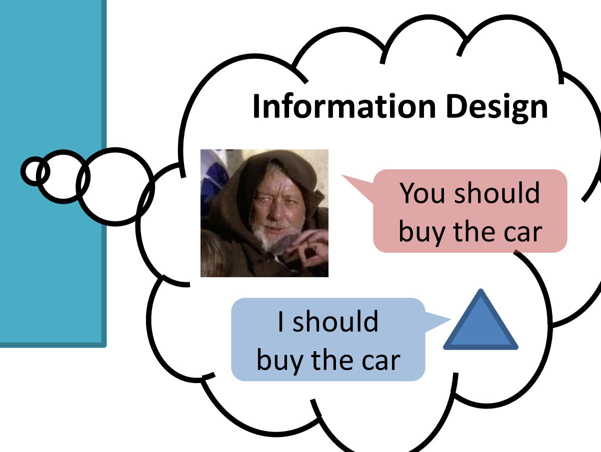

Step 1: Revelation Principle

Original equilibrium (indirect) →	**Direct equilibrium**
Deviation against original eqm ←	**Deviation against direct eqm**

Step 1: Revelation Principle

1. The players' messages to the mediator are reports of private information. In equilibrium, players always send their true private information.

Original equilibrium (indirect) | **Direct equilibrium**

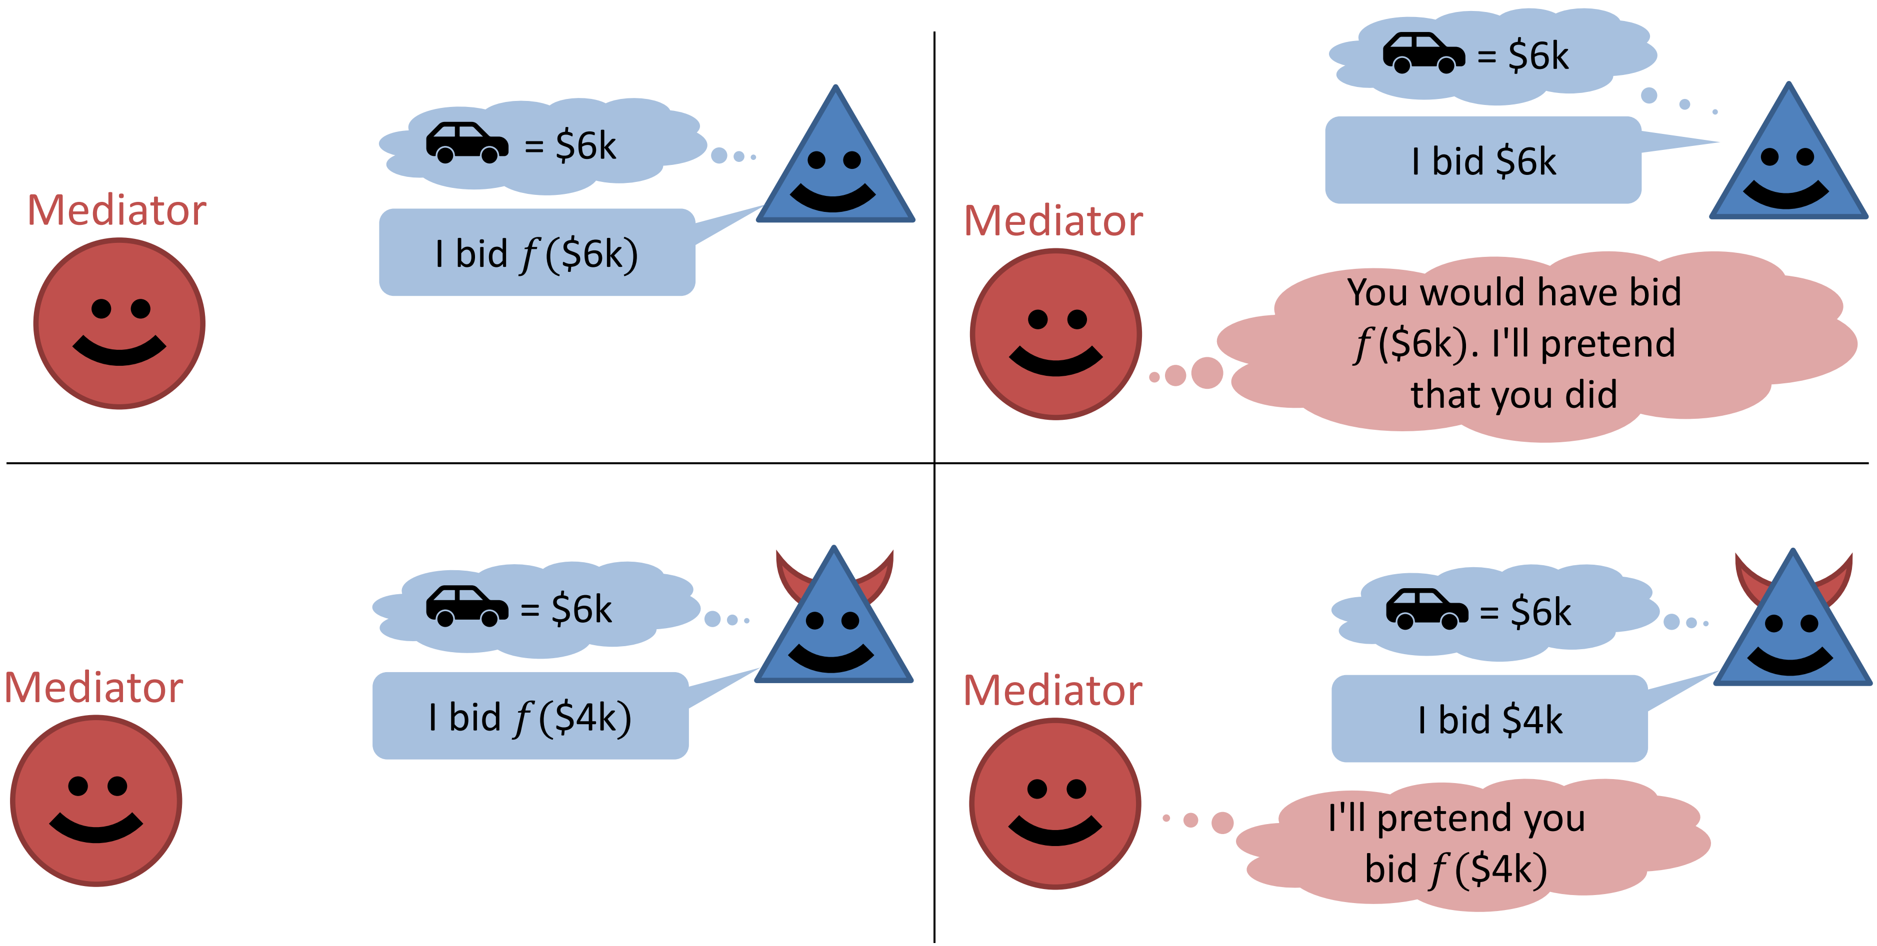

Step 1: Revelation Principle

2. The mediator's messages to the players are action recommendations. In equilibrium, players play the actions that they are recommended.

Original equilibrium (indirect) | **Direct equilibrium**

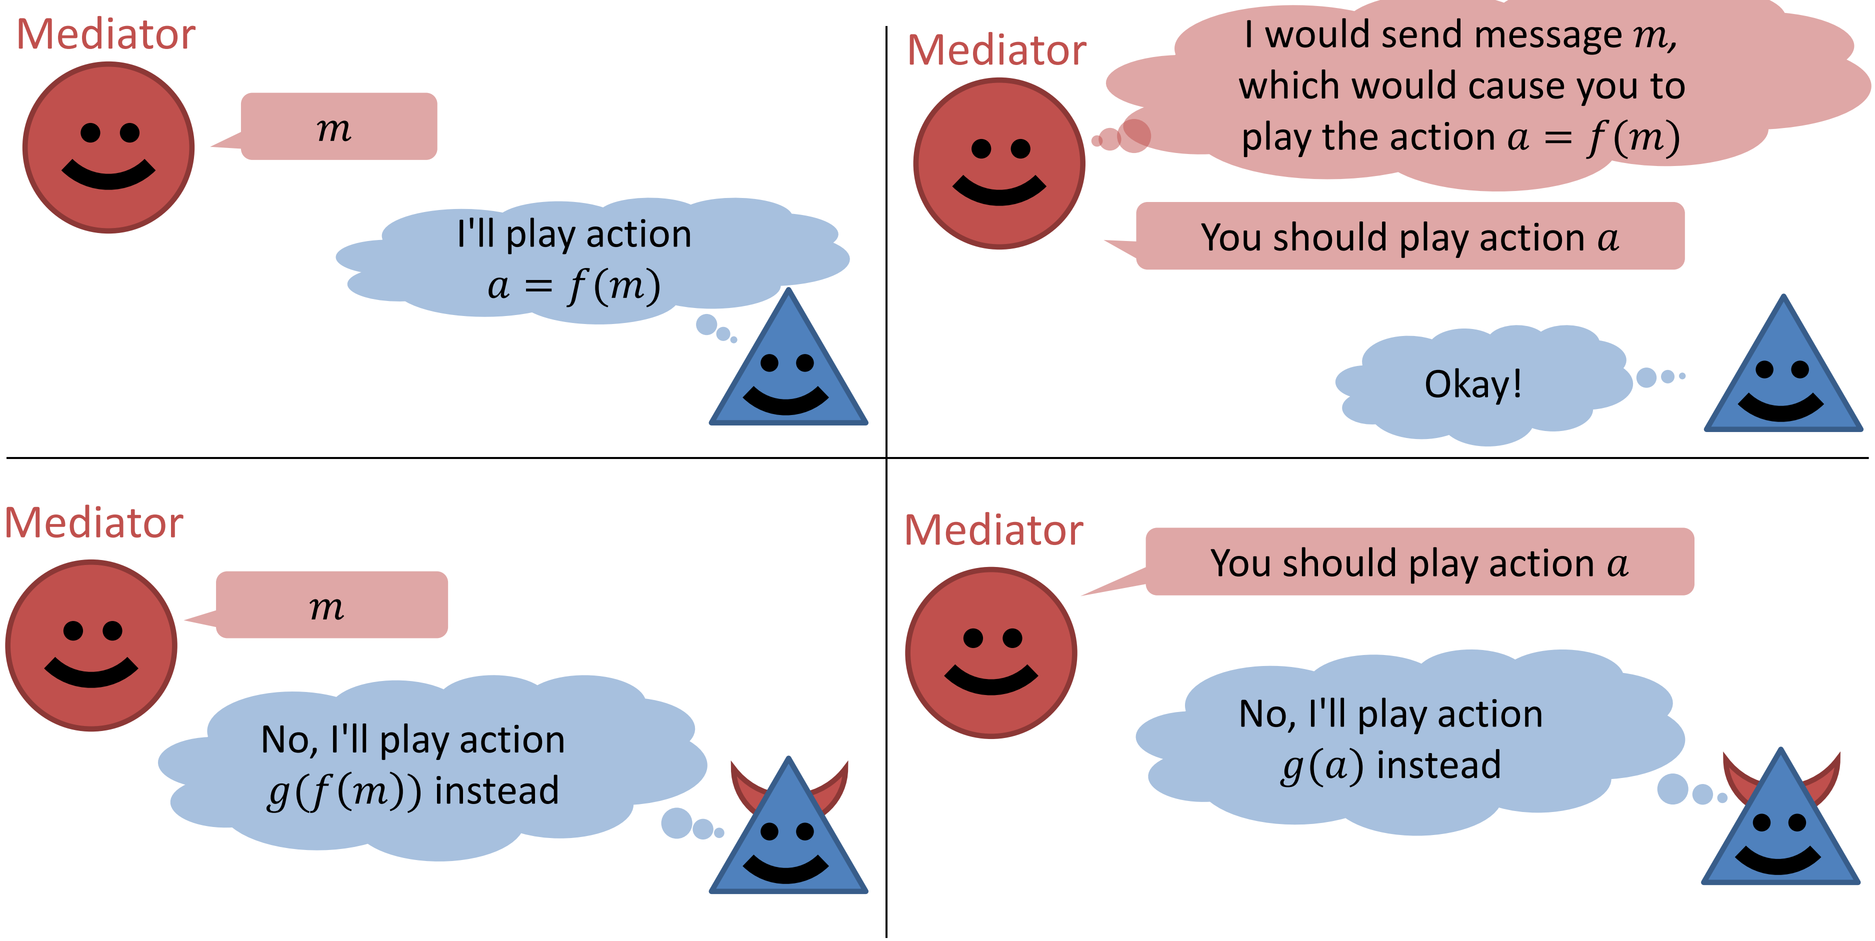

Proof in Three Steps

- Step 1: Reduce game from infinite to finite.
  ("Revelation principle" introduces **structure** to the messages) ✓

- Step 2: Reduce game from finite to polynomial.
  (Using more "without loss of generality" reductions)

- Step 3: Solve game.
  (LP duality)

Step 2: Communication Game

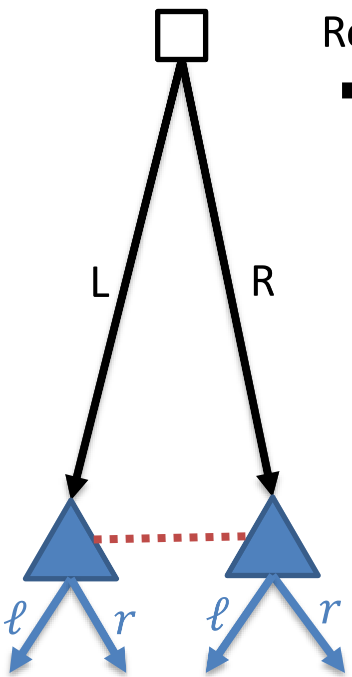

Revelation principle

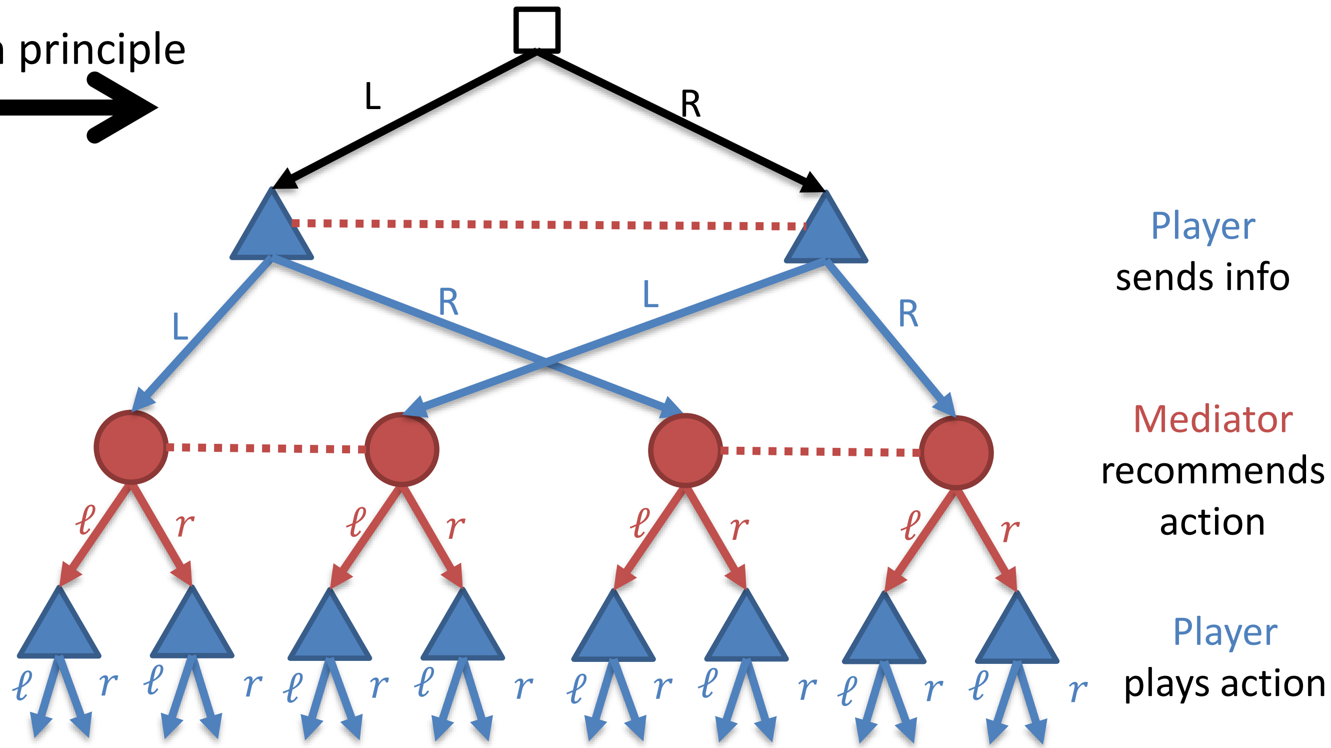

"Communication game"

Step 2: Communication Game

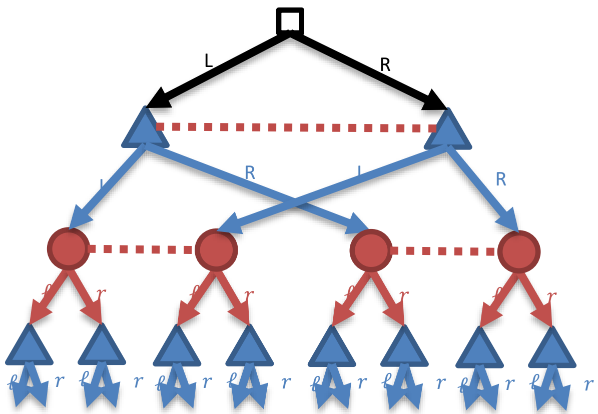

State in communication game tree

$\updownarrow$

$(s, \tau_1, \ldots, \tau_n)$

$s \in S$: state in original game tree
τ_i: transcript with player i
$n = \#$ players

#states in communication game tree
$\leq |S| \cdot (\#\text{possible messages})^{O(\text{game tree depth}) \cdot n}$

The communication game is finite! ☺
...but it is still exponentially big ☹

Observation 1: Transcripts τ_i should always correspond to some actual state s_i of the game

Proof sketch: Mediator wants to make Player i get low utility from deviating.
If τ_i doesn't correspond to an actual state,
Mediator knows that Player i deviated.
$\Rightarrow$ Mediator shouldn't give Player i any useful info
$\Rightarrow$ Players can't benefit from such τ_i

Step 2: Communication Game

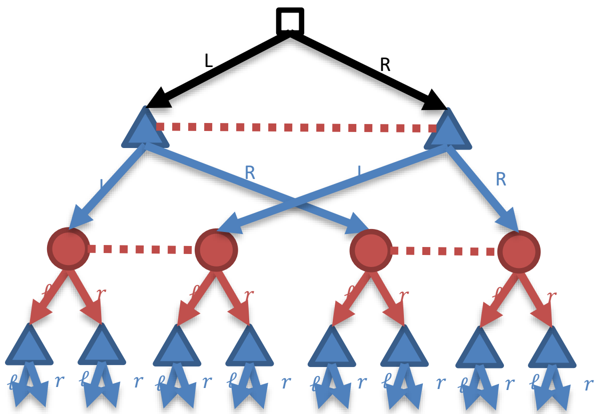

State in communication game tree

$\updownarrow$

$(s, s_1, \ldots, s_n)$

$s \in S$: state in original game tree
$s_i \in S$: state corresponding to τ_i
$n = \#$ players

#states in communication game tree $\leq |S|^{n+1}$

Much better! ☺
...but still exponential (in n) ☹

Observation 2: Only care about one deviator at a time
$\Rightarrow s_i = s$ for all but possibly one i

Step 2: Communication Game

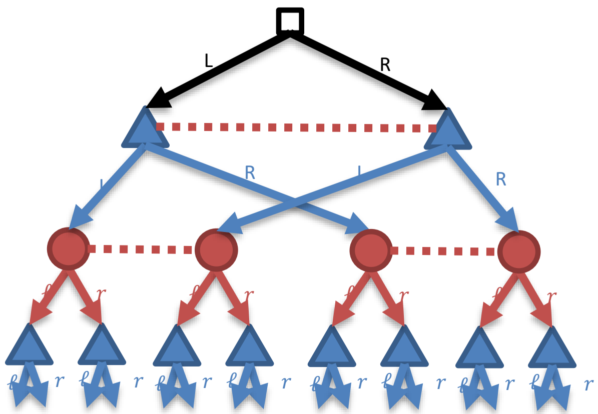

State in communication game tree

$\updownarrow$

(s, i, s_i)

$s \in S$: state in original game tree

$s_i \in S$: state corresponding to τ_i

$i \in [n] \cup \{\bot\}$: player (if any) who deviated

#states in communication game tree $\leq n \cdot |S|^2$

Yes!

Observation 2: Only care about one deviator at a time $\Rightarrow s_i = s$ for all but possibly one i

Proof in Three Steps

- Step 1: Reduce game from infinite to finite.
  ("Revelation principle" introduces **structure** to the messages)
- Step 2: Reduce game from finite to polynomial.
  (Using more "without loss of generality" reductions)
- Step 3: Solve game.
  (LP duality)

Step 3: Solving the Game

Definition:

A **communication equilibrium** is a *tuple of (possibly randomized) strategies* $(\mu, x_1, \ldots, x_n)$ such that all **players** (not incl. mediator) are best-responding:

$$u_i(\mu, x_1, \ldots, x_i, \ldots, x_n) = \max_{x_i'} u_i(\mu, x_1, \ldots, x_i', \ldots, x_n) \text{ for all } i$$

↓ Steps 1 & 2

Definition:

A **direct communication equilibrium** is a *(possibly randomized) mediator strategy μ in the communication game* such that

$$u_i(\mu, x_1^*, \ldots, x_i^*, \ldots, x_n^*) = \max_{x_i'} u_i(\mu, x_1^*, \ldots, x_i', \ldots, x_n^*) \text{ for all } i$$

where

x_i^* = **direct strategy of player i**

(Send honest info, obey recommendations)

Step 3: Solving the Game

Program:

linear in $\boldsymbol{\mu}$

$$\max u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}_1^*, \ldots, \boldsymbol{x}_i^*, \ldots, \boldsymbol{x}_n^*)$$

s.t. *constants*

linear in $\boldsymbol{\mu}$ *bilinear in* $\boldsymbol{\mu}, \boldsymbol{x}_i'$

$$u_i(\boldsymbol{\mu}, \boldsymbol{x}_1^*, \ldots, \boldsymbol{x}_i^*, \ldots, \boldsymbol{x}_n^*) \geq \max_{\boldsymbol{x}_i'} u_i(\boldsymbol{\mu}, \boldsymbol{x}_1^*, \ldots, \boldsymbol{x}_i', \ldots, \boldsymbol{x}_n^*) \text{ for all } \boldsymbol{i}$$

Definition:

A **(direct) comm eq** is a *(possibly randomized) mediator strategy* $\boldsymbol{\mu}$ *in the communication game* such that

$$u_i(\boldsymbol{\mu}, \boldsymbol{x}_1^*, \ldots, \boldsymbol{x}_i^*, \ldots, \boldsymbol{x}_n^*) = \max_{\boldsymbol{x}_i'} u_i(\boldsymbol{\mu}, \boldsymbol{x}_1^*, \ldots, \boldsymbol{x}_i', \ldots, \boldsymbol{x}_n^*) \text{ for all } \boldsymbol{i}$$

where

$\boldsymbol{x}_i^*$ = **direct strategy of player** $\boldsymbol{i}$

(Send honest info, obey recommendations)

Step 3: Solving the Game

Program:

$$\max \underbrace{\boldsymbol{c}^\top \boldsymbol{\mu}}_{\textit{linear in } \boldsymbol{\mu}}$$

s.t.

$$\underbrace{\boldsymbol{b}_i^\top \boldsymbol{\mu}}_{\textit{linear in } \boldsymbol{\mu}} \geq \max_{\boldsymbol{x}_i'} \underbrace{\boldsymbol{\mu}^\top \mathbf{A}_i \boldsymbol{x}_i'}_{\textit{bilinear in } \boldsymbol{\mu}, \boldsymbol{x}_i'} \text{ for all } \boldsymbol{i}$$

Definition:

A **(direct) comm eq** is a *(possibly randomized) mediator strategy* $\boldsymbol{\mu}$ *in the communication game* such that

$$u_i(\boldsymbol{\mu}, \boldsymbol{x}_1^*, \dots, \boldsymbol{x}_i^*, \dots, \boldsymbol{x}_n^*) = \max_{\boldsymbol{x}_i'} u_i(\boldsymbol{\mu}, \boldsymbol{x}_1^*, \dots, \boldsymbol{x}_i', \dots, \boldsymbol{x}_n^*) \text{ for all } \boldsymbol{i}$$

where

$\boldsymbol{x}_i^*$ = **direct strategy of player $\boldsymbol{i}$**

(Send honest info, obey recommendations)

Step 3: Solving the Game

Program:

$$\max_{\boldsymbol{\mu} \in \boldsymbol{X}_{\mathbf{M}}} \boldsymbol{c}^\top \boldsymbol{\mu} \quad \text{s.t.}$$

$$\boldsymbol{b}_{\boldsymbol{i}}^\top \boldsymbol{\mu} \geq \max_{\boldsymbol{x}_{\boldsymbol{i}}' \in \boldsymbol{X}_{\boldsymbol{i}}} \boldsymbol{\mu}^\top \mathbf{A}_{\boldsymbol{i}} \boldsymbol{x}_{\boldsymbol{i}}' \text{ for all } \boldsymbol{i}$$

Take duals of inner maximizations

Let $\boldsymbol{x}_{\boldsymbol{i}}' \in \boldsymbol{X}_{\boldsymbol{i}} = \{\boldsymbol{x} : \mathbf{F}_{\boldsymbol{i}} \boldsymbol{x} = \boldsymbol{f}_{\boldsymbol{i}}, \boldsymbol{x} \geq \mathbf{0}\}$

Linear program:

$$\max_{\boldsymbol{\mu} \in \boldsymbol{X}_{\mathbf{M}},\, \boldsymbol{v}_{\boldsymbol{i}} : \boldsymbol{i} \in [\boldsymbol{n}]} \boldsymbol{c}^\top \boldsymbol{\mu} \quad \text{s.t.}$$

$$\boldsymbol{b}_i^\top \boldsymbol{\mu} \geq \boldsymbol{f}_i^\top \boldsymbol{v}_i, \quad \mathbf{F}_i^\top \boldsymbol{v}_i \geq \mathbf{A}_i^\top \boldsymbol{\mu} \quad \text{for all } \boldsymbol{i}$$

□

Recap

Main theorem:

There exists **poly(size of game tree)** algorithm that computes a communication equilibrium $(\boldsymbol{\mu}, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ maximizing mediator's objective

$$u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$$

↓

Polytime algorithms for:

- Optimal sequential mechanism design
- Optimal sequential information design

...and more!

- Optimal "certification equilibria" [Forges & Koessler, *J Math Econ*'05]
- Optimal "mediated equilibria" [Monderer & Tennenholtz, *AI*'09]

Experiments: Payoff Space Plots

Medium-sized bargaining game
(≈1000 states)

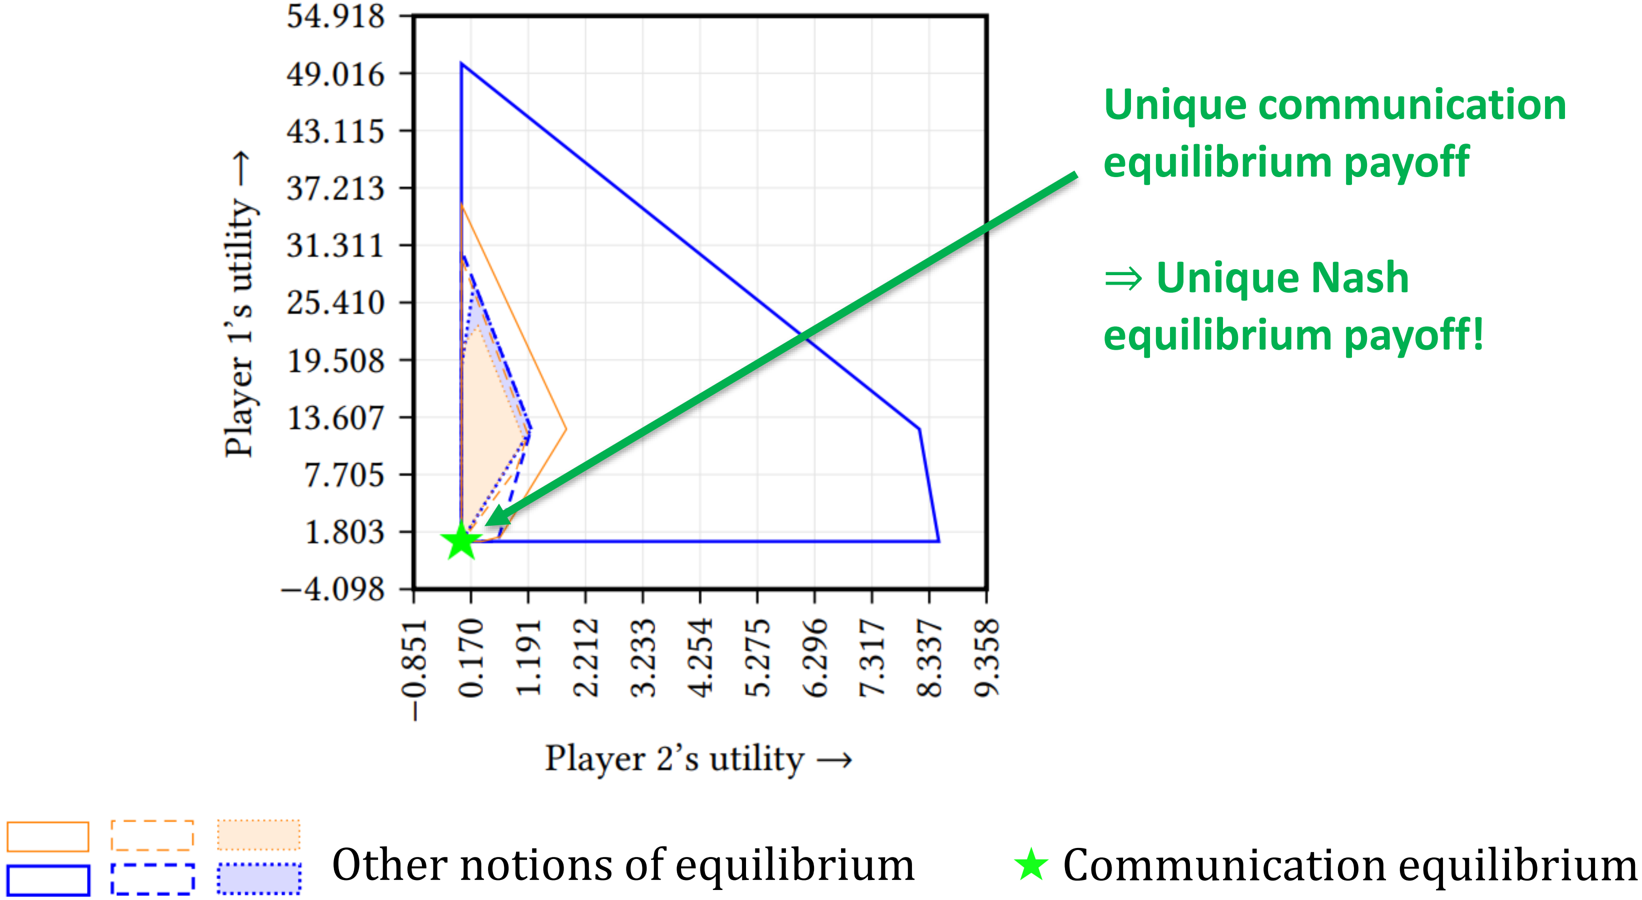

Unique communication equilibrium payoff

⇒ Unique Nash equilibrium payoff!

A Lagrangian-Based Method

find optimal mediator strategy $\boldsymbol{\mu}$

s.t. for all players i

direct strategy is a best response to $\boldsymbol{\mu}$ if all other players are direct

$$\max_{\boldsymbol{\mu}} u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}^*)$$

s.t. for all players i

$$\max_{\boldsymbol{x}_i} u_i(\boldsymbol{\mu}, \boldsymbol{x}_i, \boldsymbol{x}_{-\boldsymbol{i}}^*) \le u_i(\boldsymbol{\mu}, \boldsymbol{x}_i^*, \boldsymbol{x}_{-\boldsymbol{i}}^*)$$

A Lagrangian-Based Method

find optimal mediator strategy $\boldsymbol{\mu}$

s.t. for all players i

direct strategy is a best response to $\boldsymbol{\mu}$ if all other players are direct

$$\max_{\boldsymbol{\mu}} u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}^*)$$

s.t. for all players i

$$\max_{\boldsymbol{x}_i} u_i(\boldsymbol{\mu}, \boldsymbol{x}_i, \boldsymbol{x}^*_{-i}) \le u_i(\boldsymbol{\mu}, \boldsymbol{x}^*_i, \boldsymbol{x}^*_{-i})$$

$$\max_{\boldsymbol{\mu}} \min_{i, \boldsymbol{x}_i} u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}^*) - \lambda \cdot [u_i(\boldsymbol{\mu}, \boldsymbol{x}_i, \boldsymbol{x}^*_{-i}) - u_i(\boldsymbol{\mu}, \boldsymbol{x}^*_i, \boldsymbol{x}^*_{-i})]$$

This is a zero-sum game!

Proposition: There exists $\lambda^* > 0$ s.t. for all $\lambda > \lambda^*$:

Equilibrium strategy for max-player of this zero-sum game $\equiv$ Optimal communication equilibrium of original game

The Lagrangian as an Extensive-Form Game

mediator (Max)

deviator (Min)

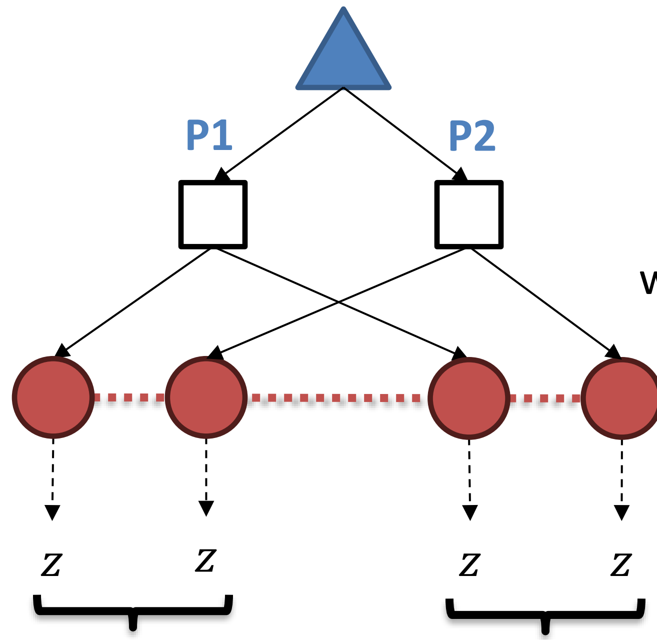

Deviator selects a player to represent

Nature selects whether **player i** can deviate (uniformly at random)

zero-sum game utility function:

$2u_{\mathrm{M}}(z) + 2\lambda u_i(z)$

Mediator plays with all **direct players**

$-2\lambda u_i(z)$

Mediator plays with all **direct players** except **player i** (controlled by **Deviator**)

$$\max_{\boldsymbol{\mu}} \min_{i, x_i} u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}^*) - \lambda \cdot [u_i(\boldsymbol{\mu}, \boldsymbol{x}_i, \boldsymbol{x}_{-i}^*) - u_i(\boldsymbol{\mu}, \boldsymbol{x}_i^*, \boldsymbol{x}_{-i}^*)]$$

The Lagrangian as an Extensive-Form Game

☺ Solving a single zero-sum game allows us to compute an optimal communication equilibrium of a multi-player game!

☹ ...but only if we knew a high-enough Lagrange multiplier λ

☹ λ depends on reward scales, so it can be quite large...

Solution #1: Set $\lambda \coloneqq 1/\varepsilon$

Theorem: Hiding game-dependent factors...

- CFR converges in averages after $1/\varepsilon^4$ iterations
- OMWU converges in averages after $1/\varepsilon^2$ iterations
- OMWU converges in iterates after $1/\varepsilon^4$ iterations

$$\max_{\boldsymbol{\mu}} \min_{i, \boldsymbol{x}_i} u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}^*) - \lambda \cdot [u_i(\boldsymbol{\mu}, \boldsymbol{x}_i, \boldsymbol{x}^*_{-i}) - u_i(\boldsymbol{\mu}, \boldsymbol{x}^*_i, \boldsymbol{x}^*_{-i})]$$

Solution #2: An Alternative Lagrangian

$$\max_{\mu} u_M(\mu, x^*) \quad \text{s.t.}$$
for all players i
$$\max_{x_i} u_i(\mu, x_i, x^*_{-i}) \le u_i(\mu, x_i^*, x^*_{-i})$$

→

find μ s.t.
$$u_M(\mu, x^*) \ge \tau$$
and for all players i
$$\max_{x_i} u_i(\mu, x_i, x^*_{-i}) \le u_i(\mu, x_i^*, x^*_{-i})$$

Algorithm: binary search

Run binary search to find $\tau \in [0,1]$. Repeat for $\log(1/\varepsilon)$ rounds:

Run an algorithm to solve the Lagrangian until either:

- it finds μ guaranteeing value $> -\varepsilon$ (branch high), or
- it proves value < 0 (branch low)

Lagrangian value 0 iff exists equilibrium μ of value $\ge \tau$

$$\max_{\mu} \min_{i, x_i} u_M(\mu, x^*) - \lambda \cdot [u_i(\mu, x_i, x^*_{-i}) - u_i(\mu, x_i^*, x^*_{-i})]$$

Solution #2: An Alternative Lagrangian

Theorem:
The last $\boldsymbol{\mu}$ found by the binary search algorithm is an ε-equilibrium whose mediator objective is at least $v^* - O(\varepsilon)$
(where v^* = optimal equilibrium mediator objective)

Algorithm: binary search
Run binary search to find $\tau \in [0,1]$. Repeat for $\log(1/\varepsilon)$ rounds:
Run an algorithm to solve the Lagrangian until either:
- it finds $\boldsymbol{\mu}$ guaranteeing value $> -\varepsilon$ (branch high), or
- it proves value < 0 (branch low)

Lagrangian value 0 iff exists equilibrium $\boldsymbol{\mu}$ of value $\geq \tau$

$$\max_{\boldsymbol{\mu}} \min \left\{ u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}^*) - \tau, -\max_{i, \boldsymbol{x}_i} [u_i(\boldsymbol{\mu}, \boldsymbol{x}_i, \boldsymbol{x}^*_{-i}) - u_i(\boldsymbol{\mu}, \boldsymbol{x}^*_i, \boldsymbol{x}^*_{-i})] \right\}$$

The Alternative Lagrangian as an Extensive-Form Game

Deviator selects a player to represent (or none)

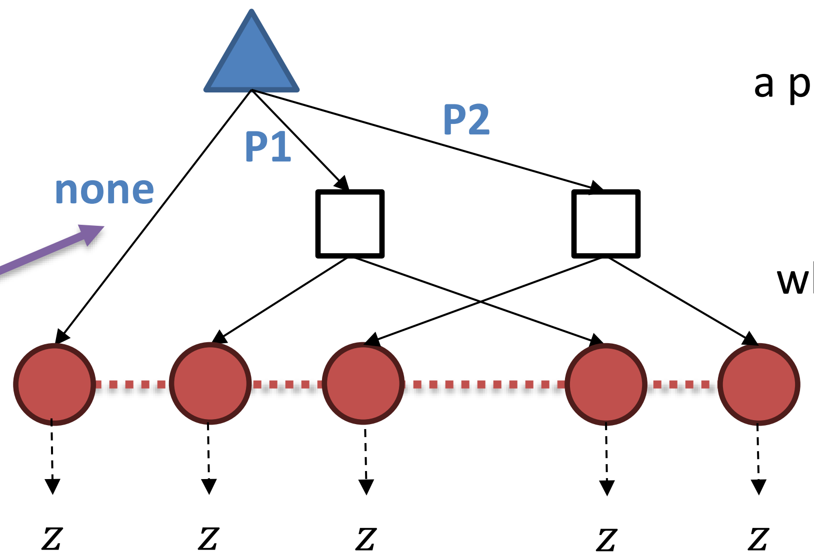

Deviator's mixed strategy probability for this edge is ~$1/\lambda$

Nature selects whether **player i** can deviate (uniformly at random)

zero-sum game utility function: $u_{\mathrm{M}}(z) - \tau$ $\quad 2u_i(z)$ $\quad -2u_i(z)$

Mediator plays with all **direct players**

Mediator plays with all **direct players** except **player i** (controlled by **Deviator**)

$$\max_{\mu} \min \left\{ u_{\mathrm{M}}(\mu, x^*) - \tau, -\max_{i, x_i} [u_i(\mu, x_i, x_{-i}^*) - u_i(\mu, x_i^*, x_{-i}^*)] \right\}$$

Which is Better?

"Direct" Lagrangian

☺ Can be formulated as an extensive-form zero-sum game

"Binary Search" Lagrangian

☺ Can be formulated as an extensive-form zero-sum game

★ If you can solve zero-sum games, you can compute optimal equilibria in various notions, optimal mechanisms, etc!

Amenable to deep RL!

"Direct" Lagrangian:

☺ Need to solve one game

☺ Last-iterate convergence is possible

☹ $O(1/\varepsilon^2)$ convergence rate (with OMWU)

☹ Extensive-form Lagrangian game has utilities whose scale depends on λ

"Binary Search" Lagrangian:

☹ Need to solve $\log(1/\varepsilon)$ games

☹ Unclear what last-iterate convergence even means

☺ $\tilde{O}(1/\varepsilon)$ convergence rate (with OMWU)

☺ Extensive-form Lagrangian game has utilities bounded by absolute constant

This really matters in practice:
deep learning solvers aren't really good at high precision!

Experiments in the Tabular Setting (not deep RL): Learning scales better than LP!

Game	# Nodes	NFCCE LP	NFCCE Ours	EFCCE LP	EFCCE Ours	EFCE LP	EFCE Ours	COMM LP	COMM Ours	CERT LP	CERT Ours
B2222	1573	0.00s	0.00s	0.00s	0.01s	0.00s	0.02s	2.00s	1.49s	0.00s	0.02s
B2322	23 839	0.00s	0.01s	3.00s	0.69s	9.00s	1.60s	timeout	4m 41s	2.00s	1.24s
B2323	254 239	6.00s	0.33s	1m 21s	14.23s	3m 40s	44.87s	timeout	timeout	37.00s	40.45s
B2324	1 420 639	38.00s	2.73s	timeout	3m 1s	timeout	10m 48s	timeout	timeout	timeout	6m 14s
D32	1017	0.00s	0.01s	0.00s	0.02s	12.00s	0.40s	0.00s	0.06s	0.00s	0.01s
D33	27 622	2m 17s	12.93s	timeout	1m 46s	timeout	timeout	timeout	4m 37s	4.00s	3.14s
GL3	7735	0.00s	0.01s	1.00s	0.02s	0.00s	0.01s	timeout	7.72s	0.00s	0.02s
K35	1501	49.00s	0.76s	46.00s	0.67s	57.00s	0.55s	1.00s	0.03s	0.00s	0.01s
L3132	8917	26.00s	0.59s	8m 43s	5.13s	8m 18s	6.10s	8.00s	3.46s	1.00s	0.10s
L3133	12 688	38.00s	0.94s	20m 26s	8.88s	21m 25s	6.84s	12.00s	3.40s	1.00s	0.22s
L3151	19 981	timeout	15.12s	timeout	timeout	timeout	timeout	timeout	16.73s	2.00s	0.21s
L3223	15 659	4.00s	0.44s	1m 10s	2.94s	2m 2s	5.52s	19.00s	18.19s	1.00s	0.61s
L3523	1 299 005	timeout	1m 7s	timeout	timeout	timeout	timeout	timeout	timeout	timeout	2m 58s
S2122	705	0.00s	0.00s	0.00s	0.01s	0.00s	0.02s	2.00s	0.35s	0.00s	0.02s
S2123	4269	0.00s	0.01s	1.00s	0.06s	1.00s	0.15s	1m 33s	59.63s	1.00s	0.15s
S2133	9648	1.00s	0.02s	3.00s	0.11s	3.00s	0.49s	timeout	12m 11s	2.00s	0.92s
S2254	712 552	1m 58s	7.43s	timeout	22.01s	timeout	3m 34s	timeout	timeout	timeout	2m 42s
S2264	1 303 177	3m 43s	11.74s	timeout	39.23s	timeout	timeout	timeout	timeout	timeout	timeout
TP3	910 737	1m 38s	7.44s	timeout	13.76s	timeout	13.46s	timeout	timeout	timeout	26.70s
RS212	598	0.00s	0.00s	0.00s	0.00s	0.00s	0.00s	2.00s	0.01s	0.00s	0.00s
RS222	734	0.00s	0.00s	0.00s	0.00s	0.00s	0.00s	3.00s	0.01s	0.00s	0.00s
RS213	6274	timeout	14.68s	timeout	15.54s	timeout	23.37s	6m 25s	8.74s	0.00s	0.02s
RS223	6238	timeout	timeout	timeout	timeout	timeout	timeout	8m 54s	4.00s	1.00s	0.01s

Here, we used our “direct Lagrangian” algorithm.

Experiments on back-to-back auctions among budget-constrained bidders

2 items. Valuations of 2nd item drawn after 1st auction. Used tabular CFR+ to solve

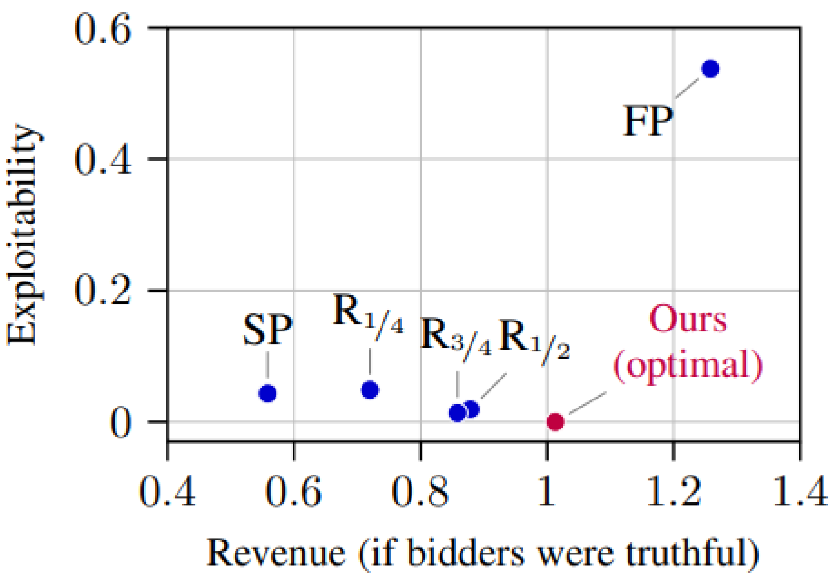

4 items. Valuations drawn after previous auction. Used deep learning to solve binary search Lagrangians!

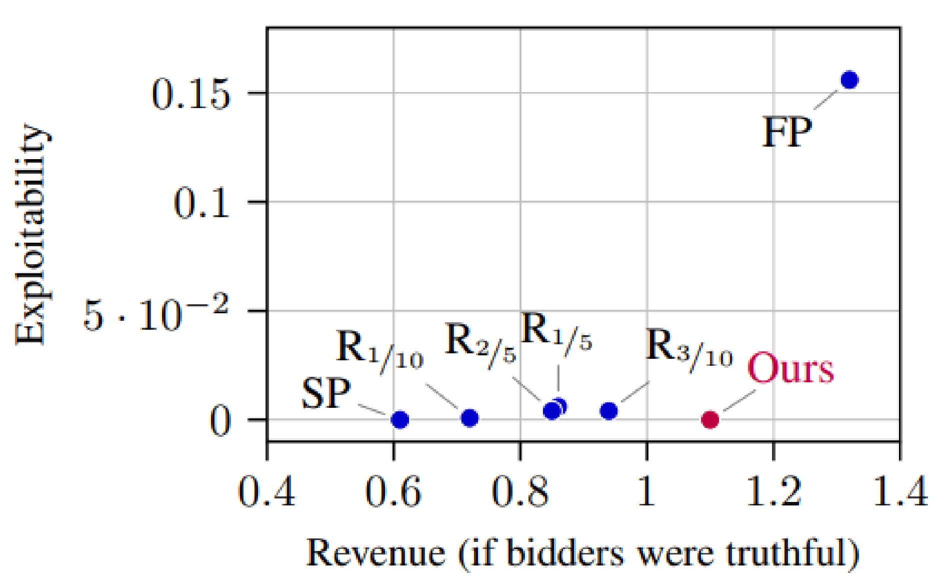

FP = first price (highly exploitable, as expected, but revenue-maximizing if bidders are truthful)
SP = second price with no reserve
R_p = second price with reserve price p

Extensive-Form Correlated Equilibria

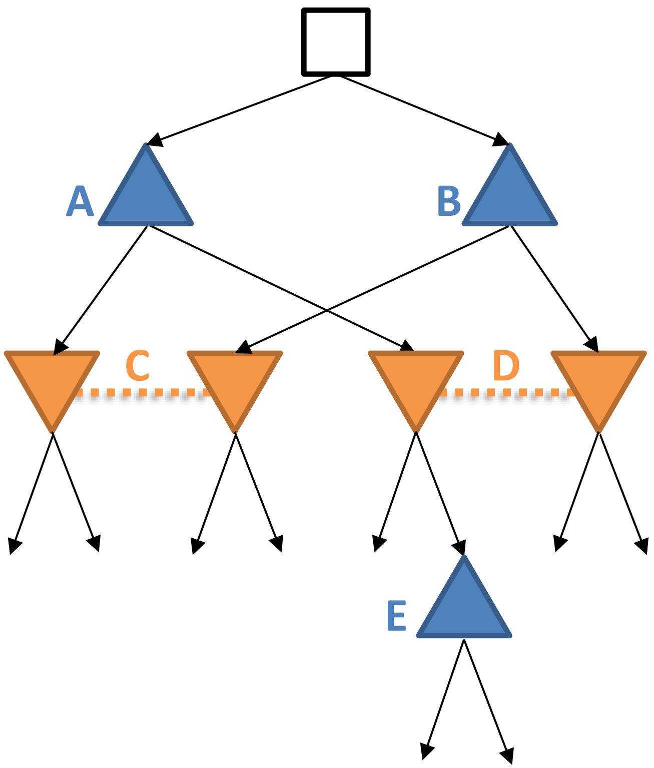

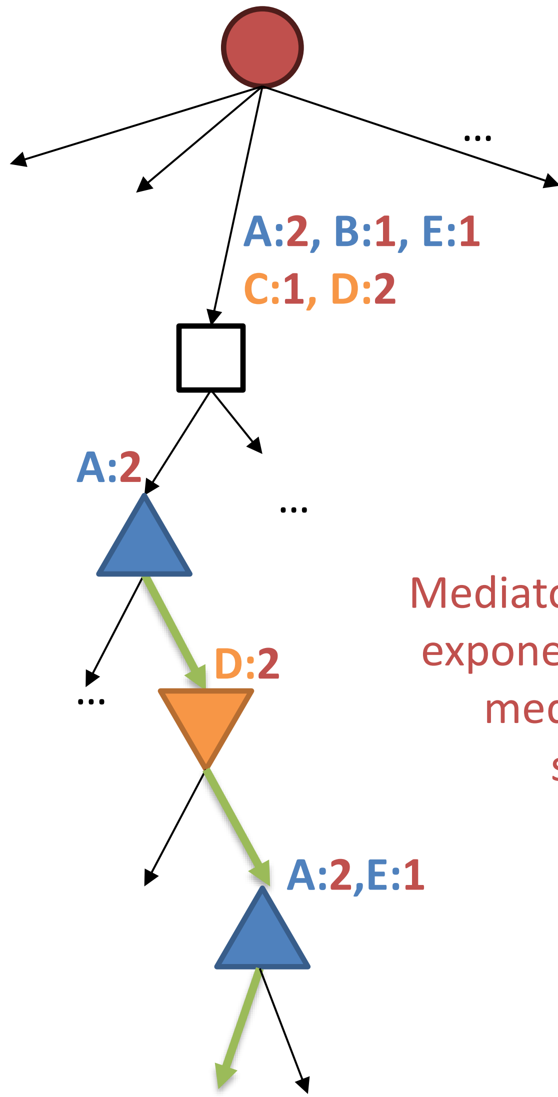

Problem:
Mediator-augmented game is exponentially large, because mediator can pick any strategy profile

Extensive-Form Correlated Equilibria

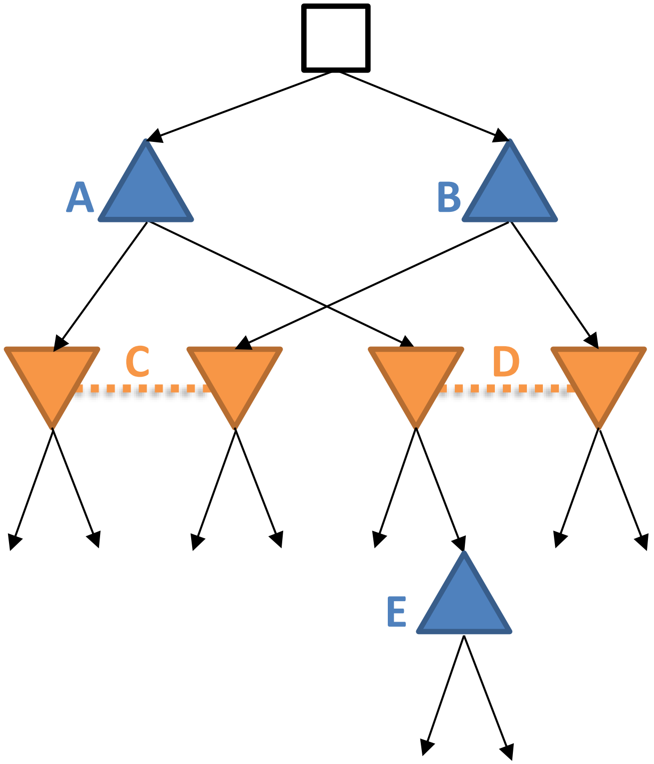

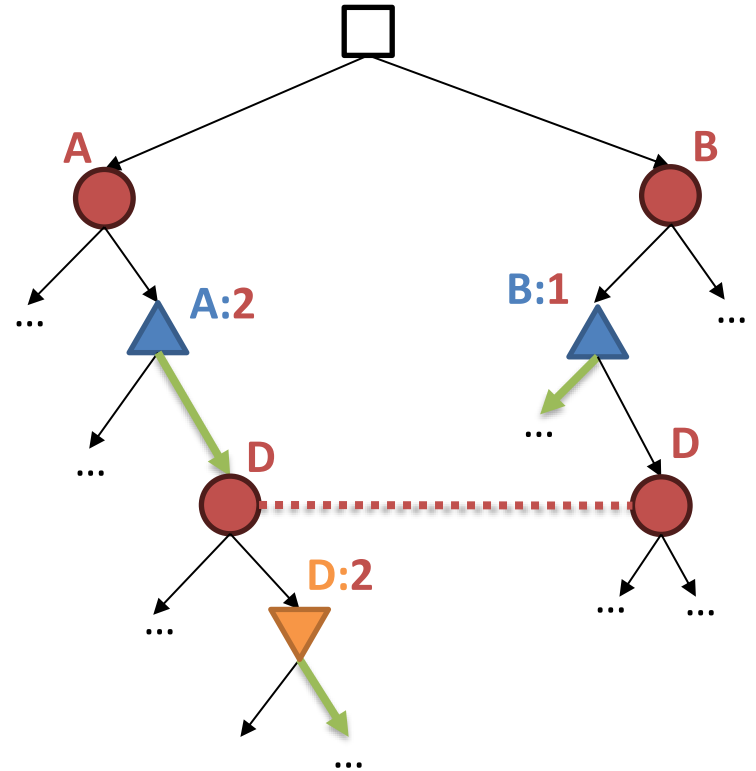

Pure strategy of mediator = Assignment of one action to each infoset of original game = Pure profile in original game

Extensive-Form Correlated Equilibria

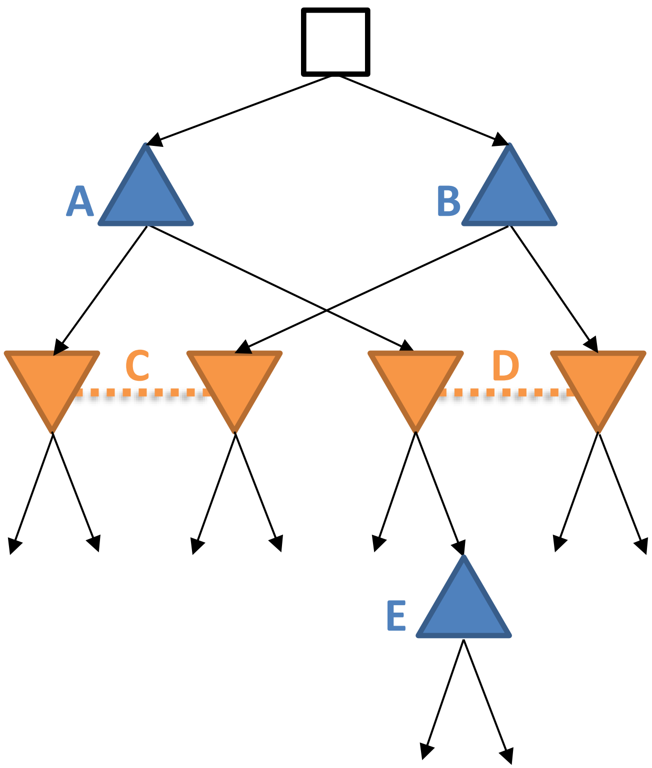

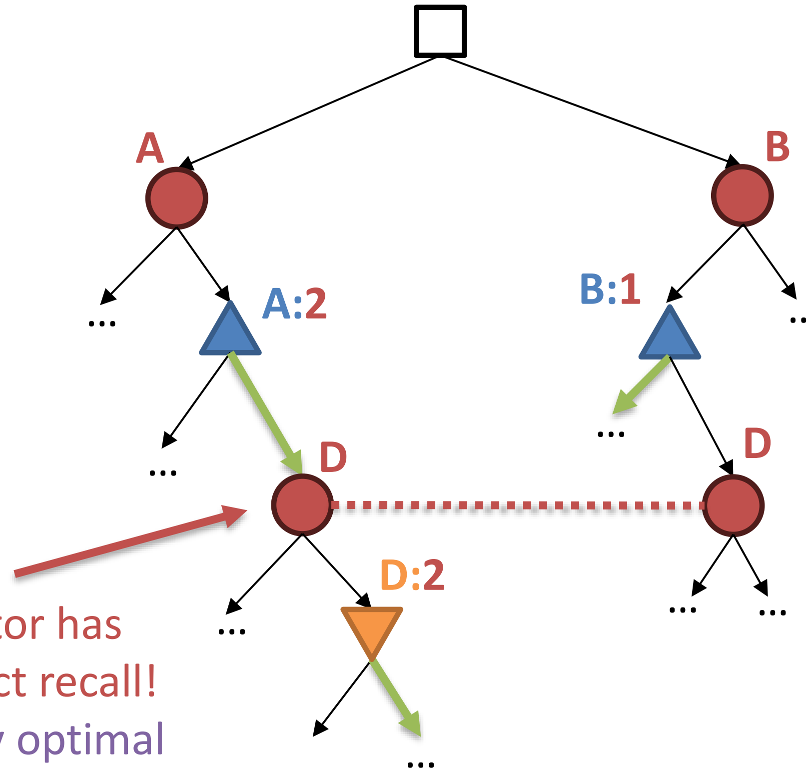

Mediator has imperfect recall!
This is why optimal EFCE is **NP-hard**

Mixed strategy of mediator **(not behavioral!)** = Distribution over pure strategy profiles in the original game = Correlated profile in original game

EFCE = **strategy for the mediator**
s.t. if the mediator **commits** to that strategy,
direct profile **is Nash equilibrium** for other players

EFCE vs Information Design vs Mechanism Design

(In the augmented game)

	Private information		Actions	
	Mediator	**Players**	**Mediator**	**Players**
Mechanism design	None	Imperfect information (e.g., private types/values for mechanism design)	Selects mechanism outcome	Type reports
Information design	Perfect information		Action recommendations	In-game actions
Extensive-form correlated equilibria	Only information of current player			

EFCE ≈ information design + privacy constraints!

Important subclasses of the general problem

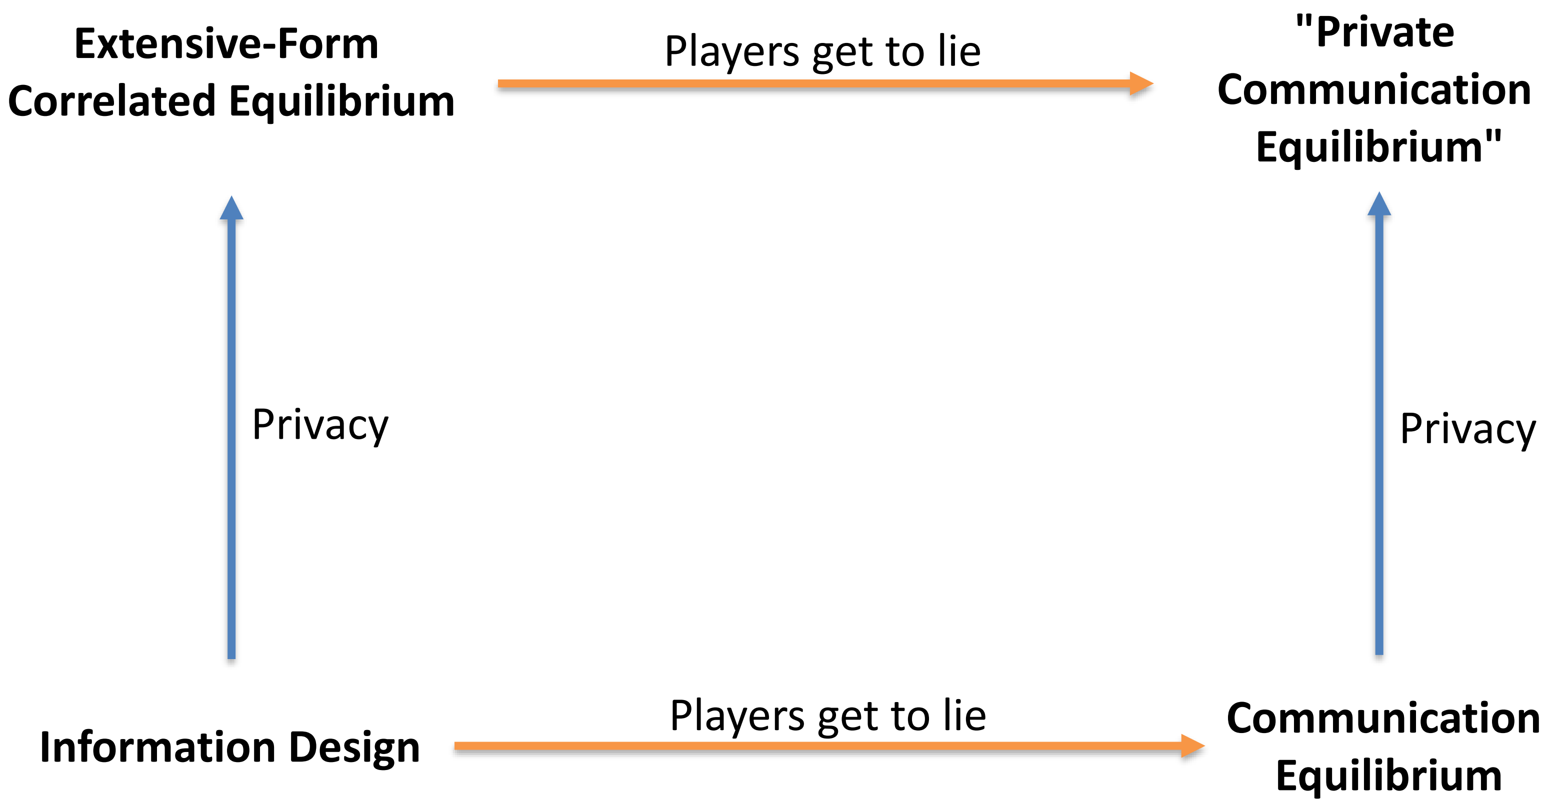

Zhang & Sandholm *NeurIPS*'22

All the existing solution concepts that we've discussed, and several more, are connected!

Normal-form correlated equilibrium [Aumann *J Math Econ*'74]
Normal-form coarse-correlated equilibrium [Moulin & Vial *Int J Game Theory*'78]
Extensive-form correlated equilibrium [von Stengel & Forges *Math of OR*'08]
Extensive-form coarse-correlated equilibrium [Farina, Bianchi, Sandholm *AAAI*'20]
Communication equilibrium [Myerson *Econometrica*'86; Forges *Econometrica*'86]
Certification equilibrium [Forges & Koessler *J Math Econ*'05]
Mediated equilibrium [Monderer & Tennenholtz *AIJ*'09]
(Sequential) mechanism design
(Sequential) information design/Bayesian persuasion
[Kamenica & Gentzkow *Am Econ Rev*'11]

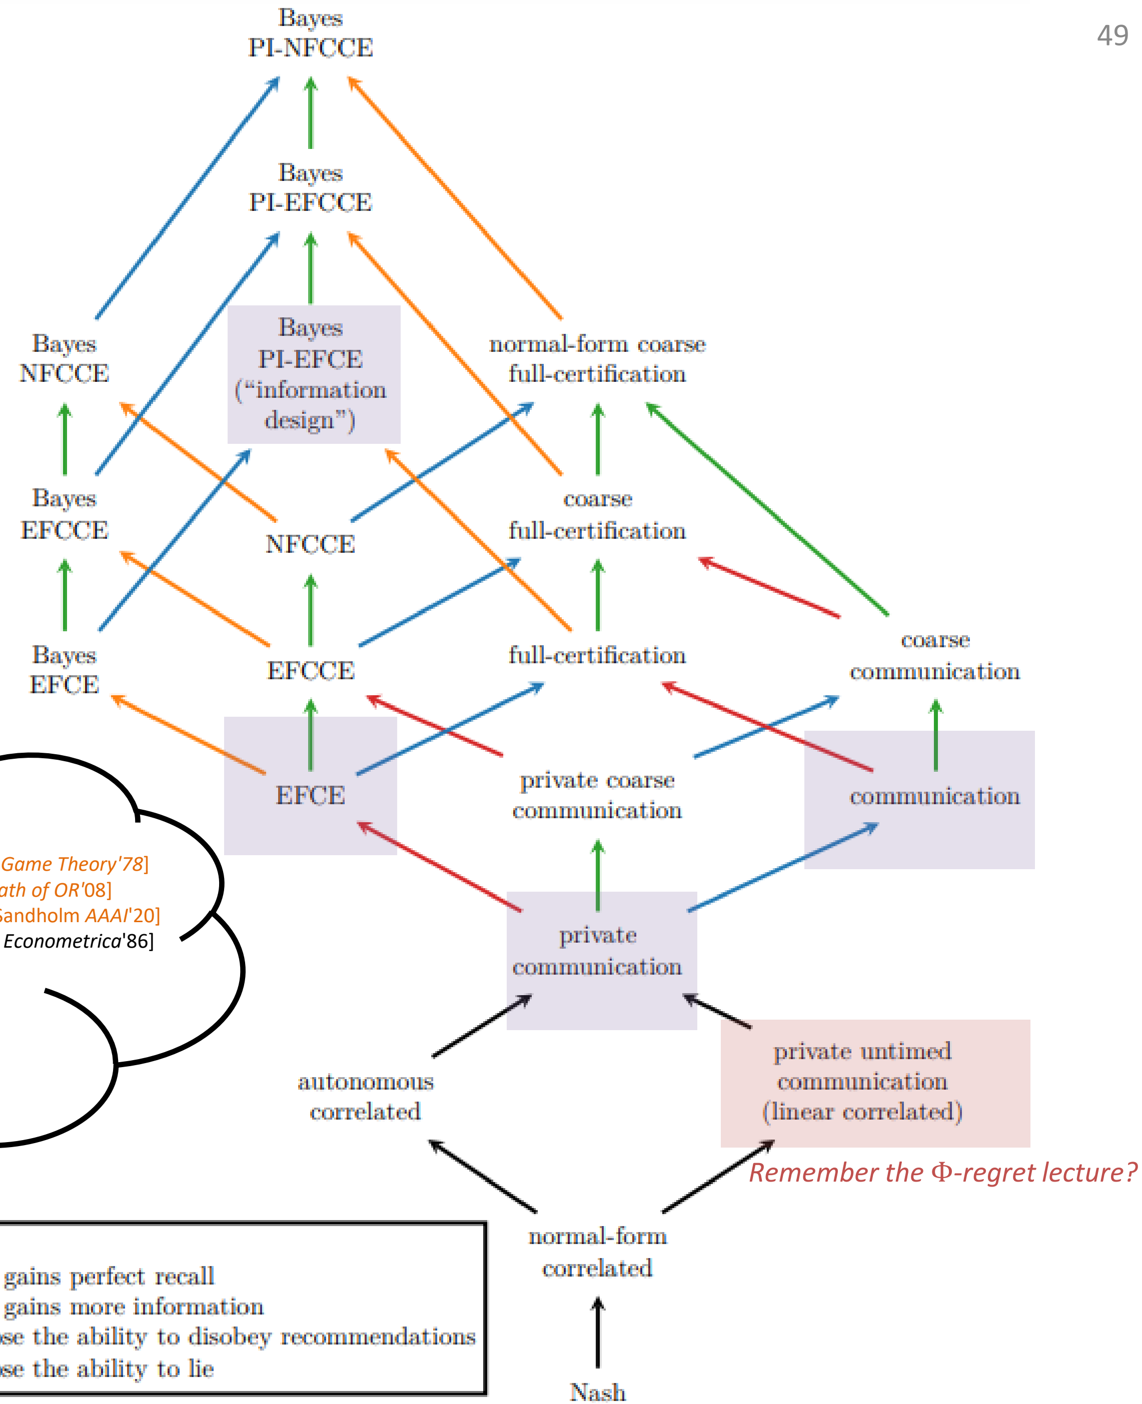

Legend:
- Mediator gains perfect recall
- Mediator gains more information
- Players lose the ability to disobey recommendations
- Players lose the ability to lie

Bibliographic Notes

Many special cases independently analyzed as separate problems.
This talk can be viewed as a **unifying framework** for these results & more!

LP-based algorithms for finding optimal equilibria:

- Mechanism design
  - Single-shot [Conitzer & Sandholm *UAI* 2002; Sandholm CP 2003]
  - One player [Zhang & Conitzer *NeurIPS* 2021]
  - Auctions [Papadimitriou, Pierrakos, Psomas, Rubinstein *GEB* 2014]
- Sequential information design
  - One player [Gan, Majumdar, Radanovic, Singla *AAAI* 2022]
  - Multiple myopic players [Wu, Zhang, Feng, Wang, Yang, Jordan *EC* 2022]
- Optimal correlated equilibria [Zhang, Farina, Celli, Sandholm *EC* 2022]

Lagrangians:

- "Direct" Lagrangian in the single-step mechanism design case
  [Dütting, Feng, Narasimhan, Parkes, Ravindranath *JACM* 2023]
- "Binary search" Lagrangian stated (but not analyzed) for EFCE and NFCCE
  [Farina, Ling, Fang, Sandholm *NeurIPS* 2019]

Bibliographic Notes

- No-regret learning algorithms for computing **one** equilibrium
  - Extensive-form correlated equilibrium
    [Farina, Celli, Marchesi, Gatti *JACM* 2022]
  - Communication equilibrium
    [Fujii *arXiv* 2023]
  - Linear correlated equilibrium
    [Farina & Pipis *NeurIPS* 2023; Zhang, Farina, Sandholm *ICLR* 2024]
  - Normal-form correlated equilibrium
    $n^{\tilde{O}(1/\varepsilon)}$ convergence rate, where *n* is the number of nodes
    [Peng & Rubinstein *arXiv* 2023; Dagan, Daskalakis, Fishelson & Golowich *arXiv* 2023]
    Open: can $\mathbf{poly}(\boldsymbol{n}, \mathbf{1}/\boldsymbol{\varepsilon})$ rate be achieved as in the other equilibrium concepts above?

- Other applications of mediators:
  - Team-correlated equilibria in adversarial team games
    [Carminati, Cacciamani, Ciccone, Gatti *ICML* 2022]
    [**Zhang**, Farina, Sandholm *ICML* 2023]
    [**Zhang** & Sandholm *AAAI* 2022]
  - Hidden-role games, such as *Avalon*
    [Carminati, **Zhang**, Farina, Gatti, Sandholm *arXiv* 2023]

- Future research: Large-scale experiments (e.g., in sequential auctions) with deep RL?