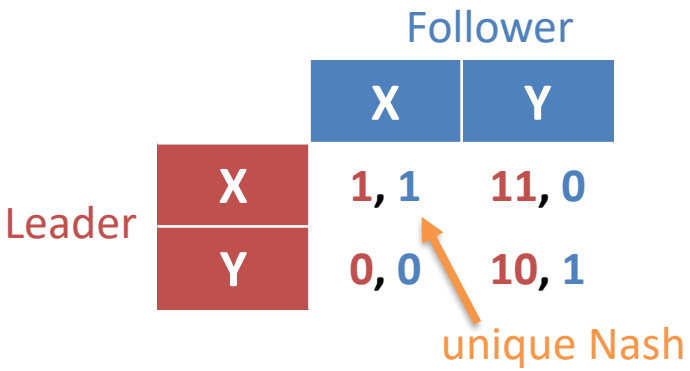
General mechanism design via extensive-form games

Brian Hu Zhang

Problem 1: Stackelberg equilibria in normal-form games



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

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

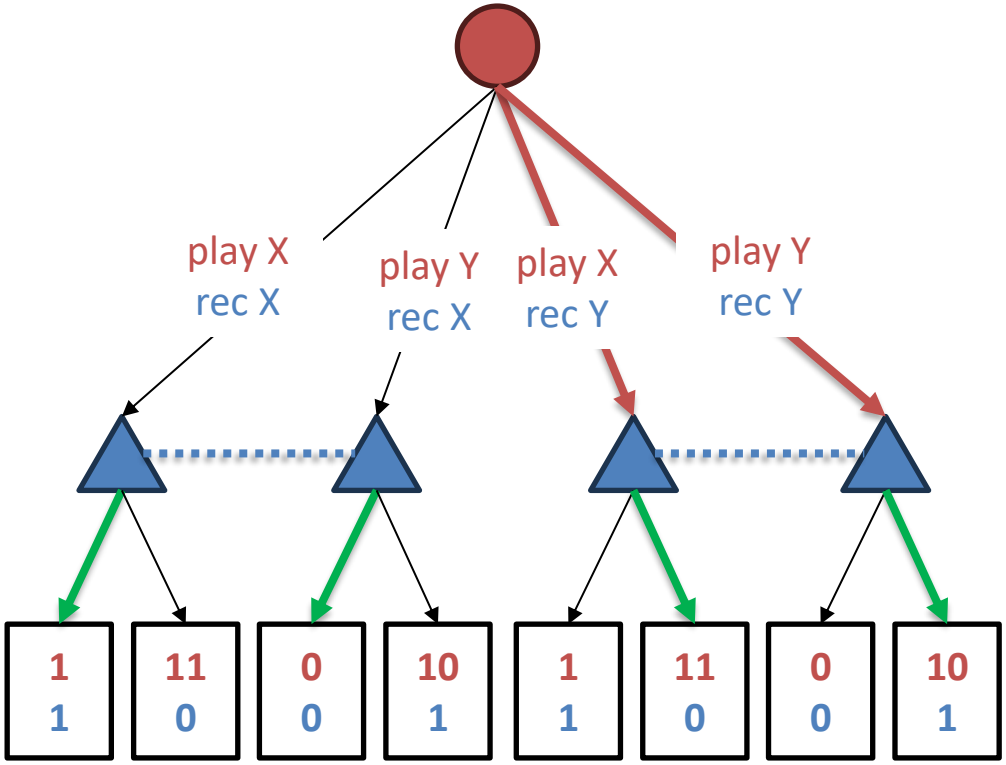
$$\text{s.t. } u_F(x_L, x_F) \geq u_F(x_L, a_F) \ \forall a_F \in A_F$$

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

(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

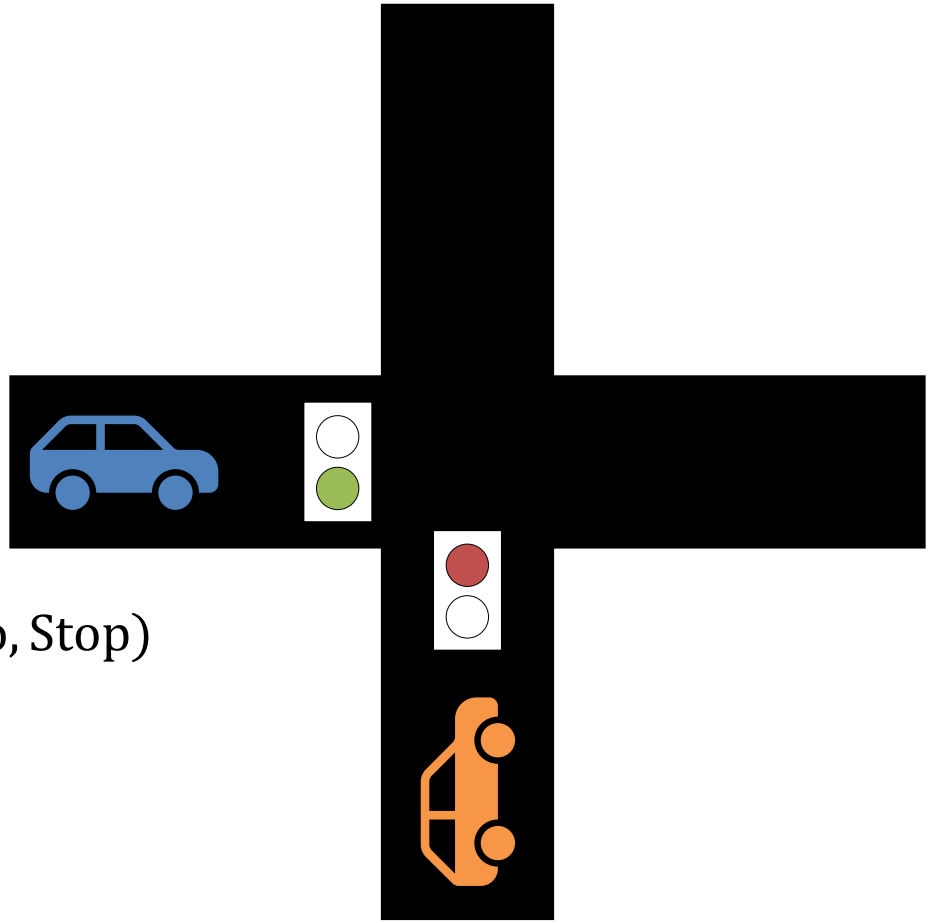


Stackelberg equilibrium = strategy for leader
s.t. holding leader's strategy fixed,
direct strategy is a **best response** for follower

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

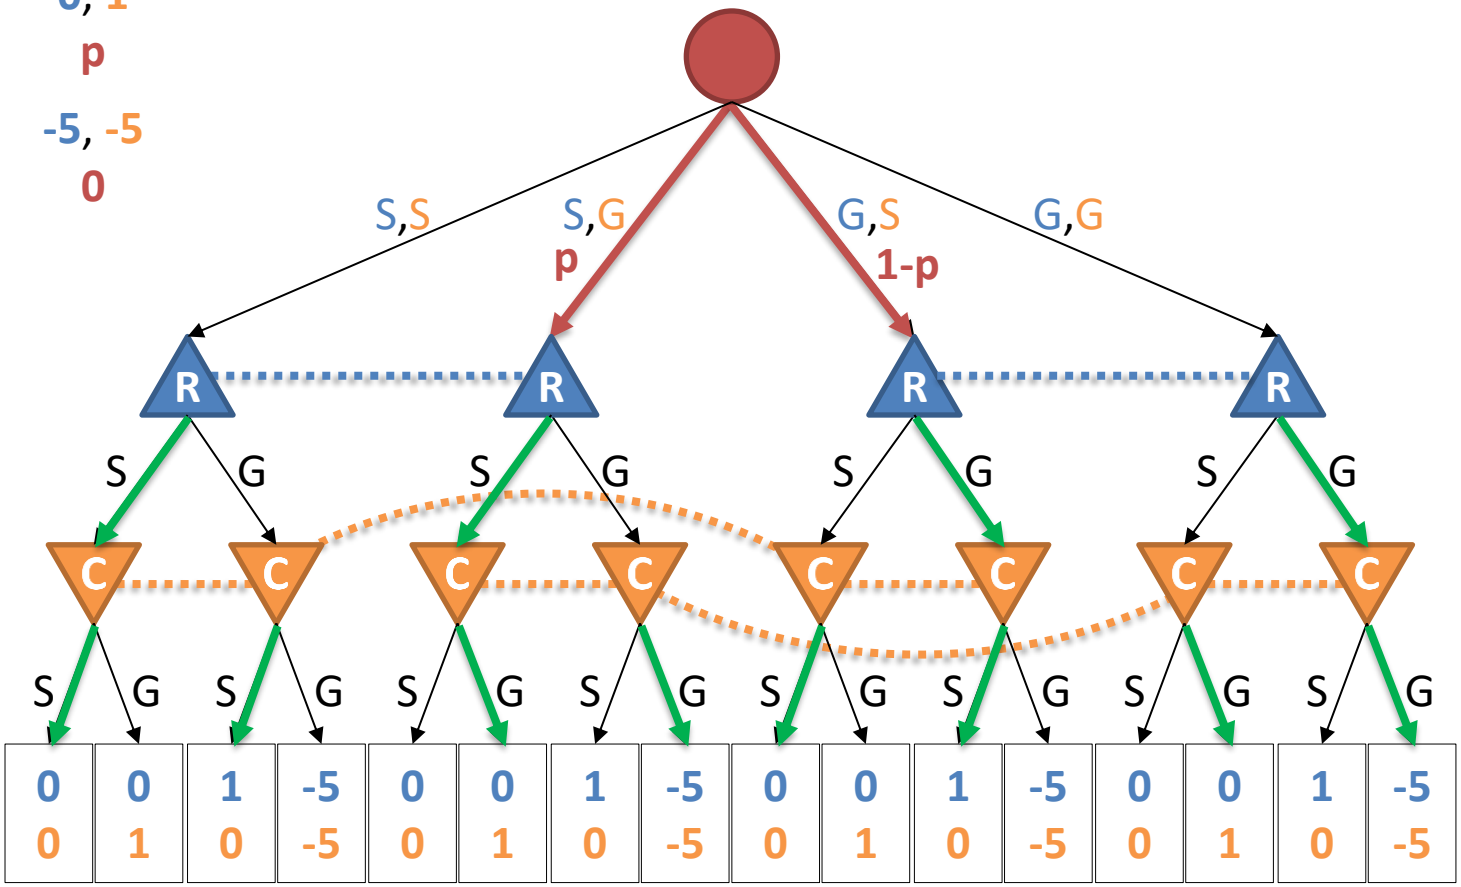
Player strategies are **direct**: incentivized to follow recommendations

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

How to maximize (expected) revenue?

Auctioneer

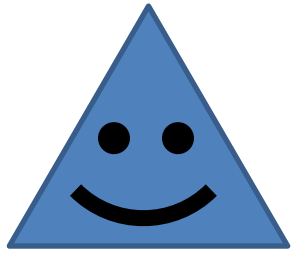
I don't know how much buyers value car

I could ask them directly, but:
- what if they lie?
- is that best?



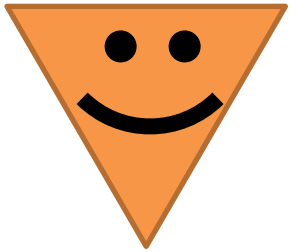
$6k

Buyer 1



$4k

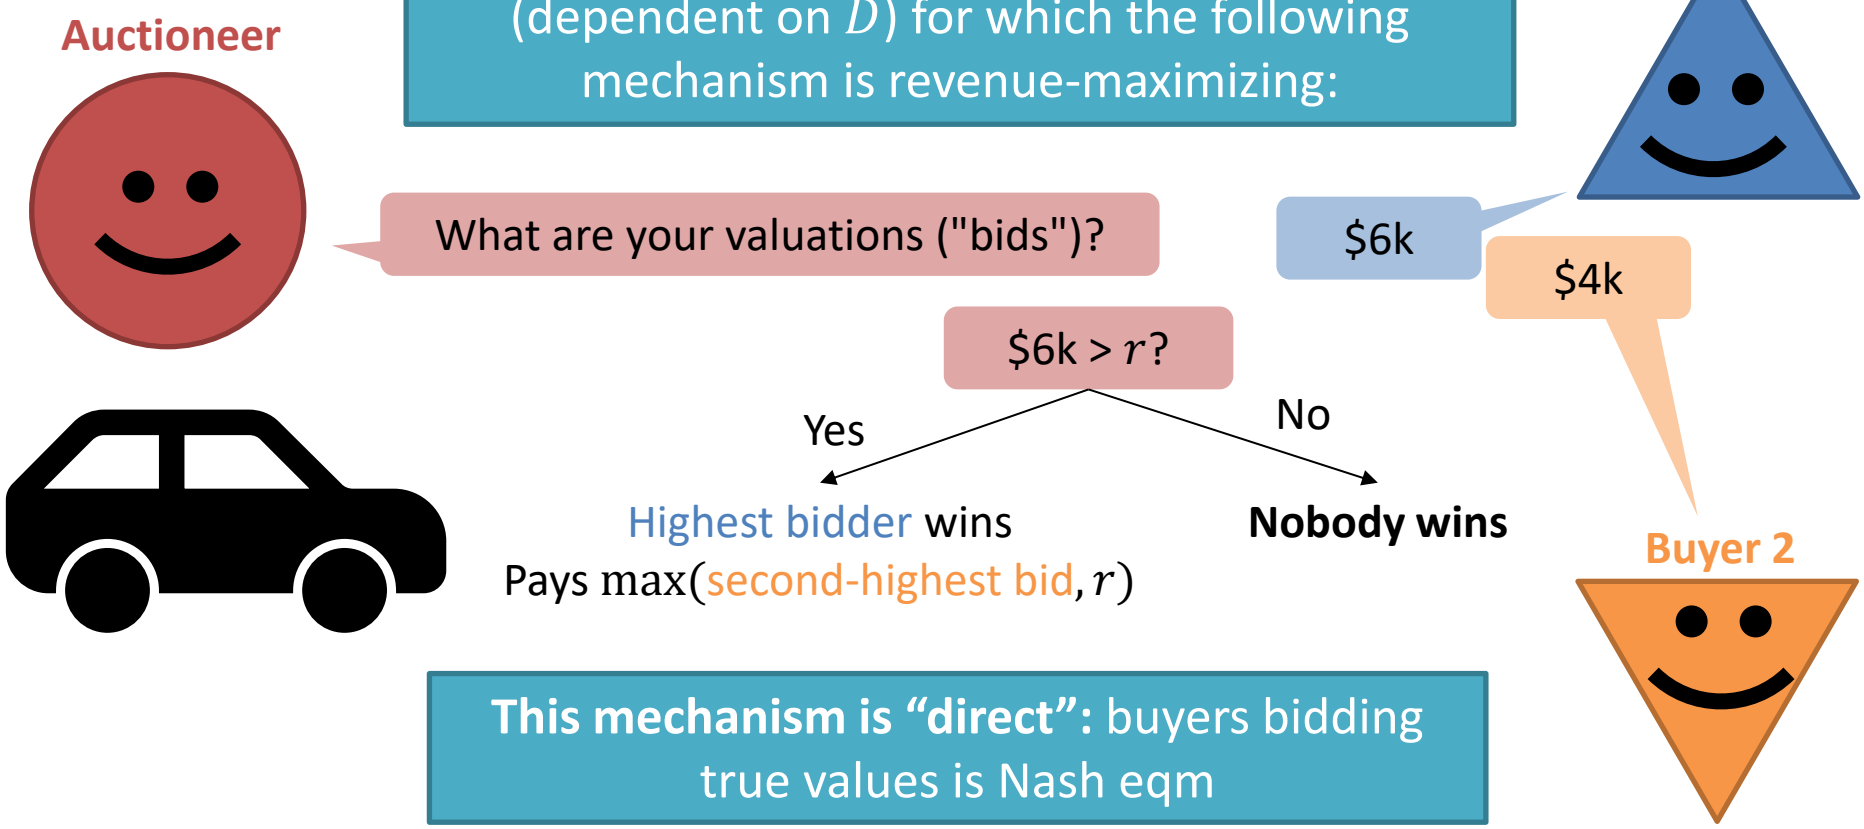
Buyer 2



Problem 3: Mechanism design

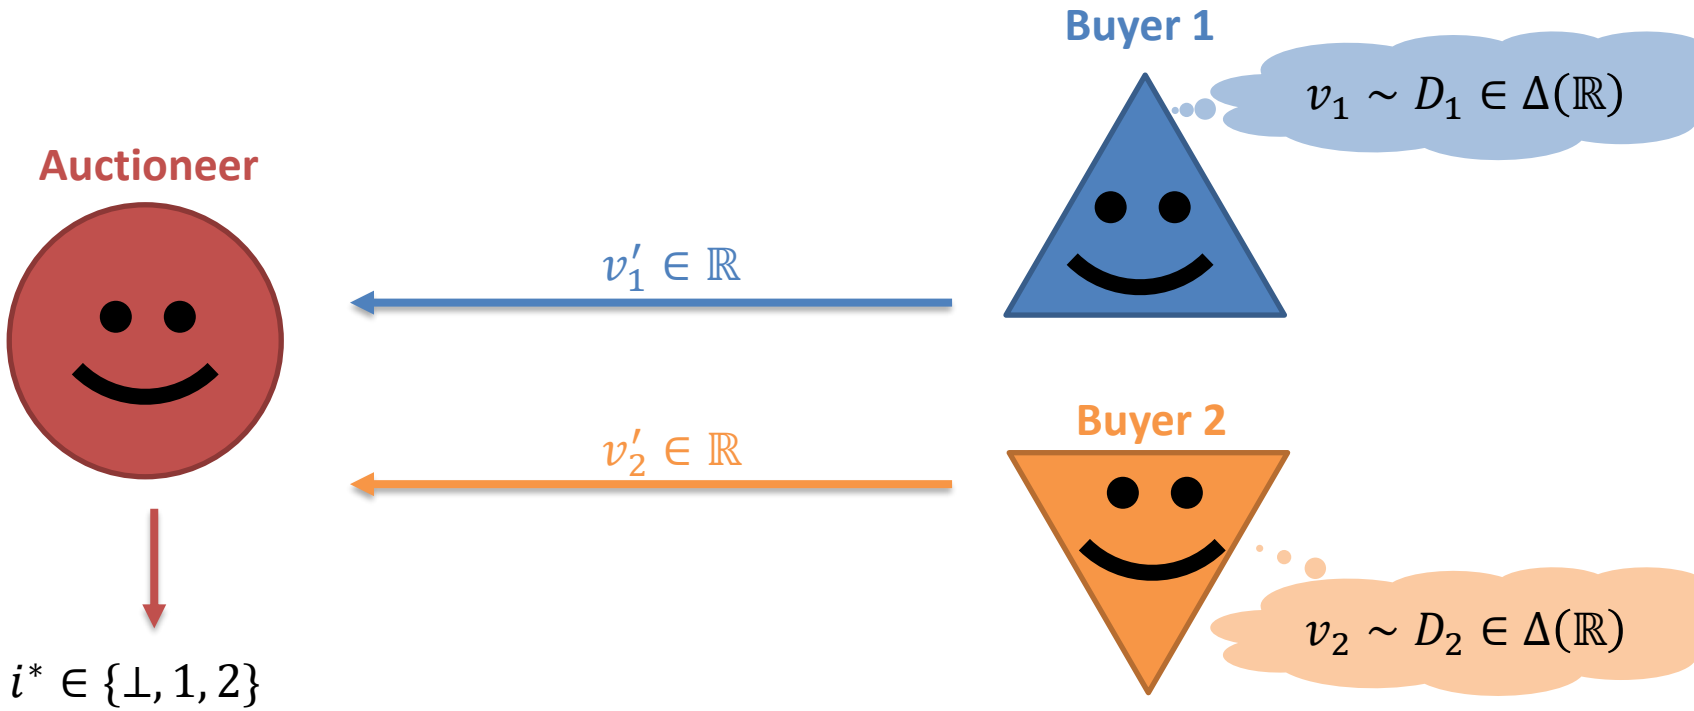
Theorem (Myerson, *Math of OR*'81):

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:



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 $v_i' = v_i$

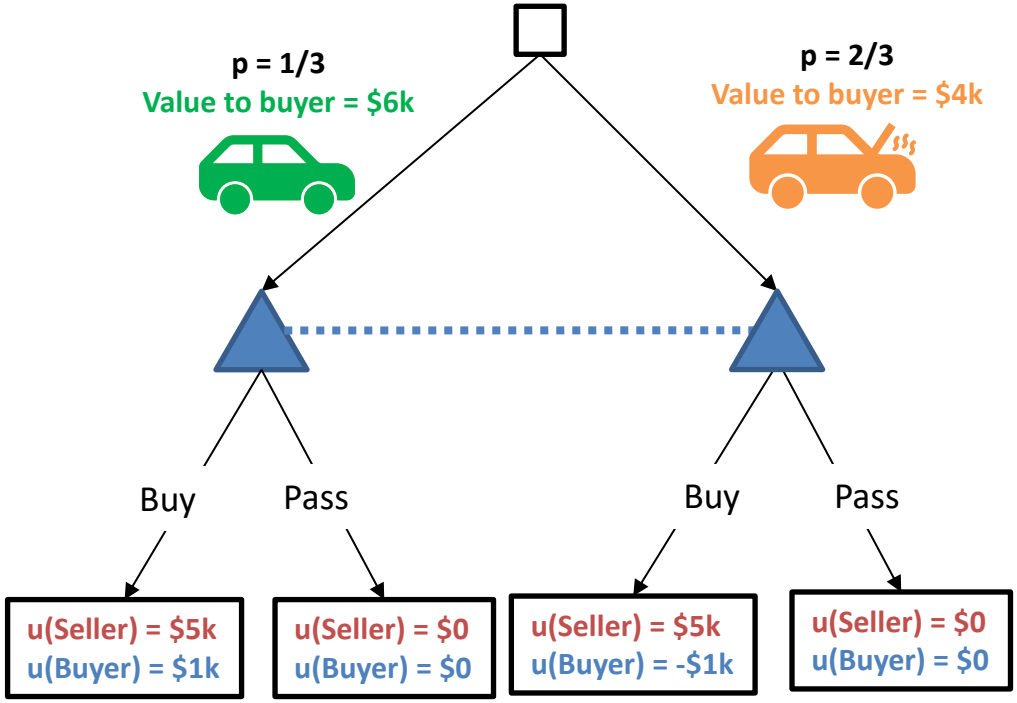
Optimal mechanism = **strategy for the auctioneer**
s.t. holding auctioneer's strategy fixed,
direct profile is Nash equilibrium for other players

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

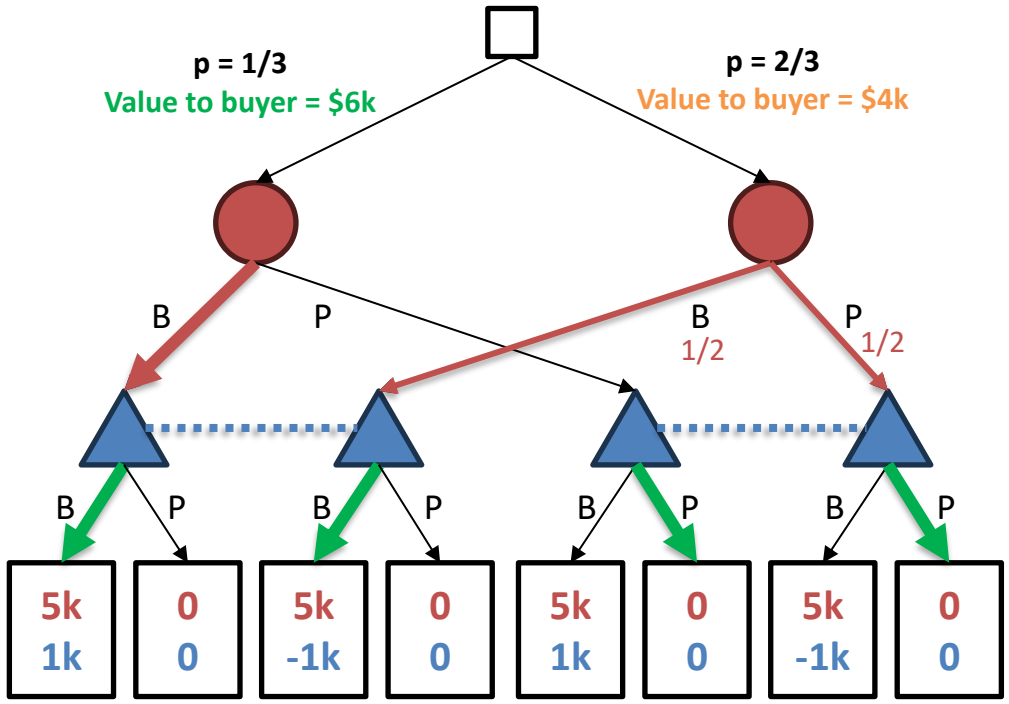
As before: The seller can commit, and send messages

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



You should buy the car

I should buy the car

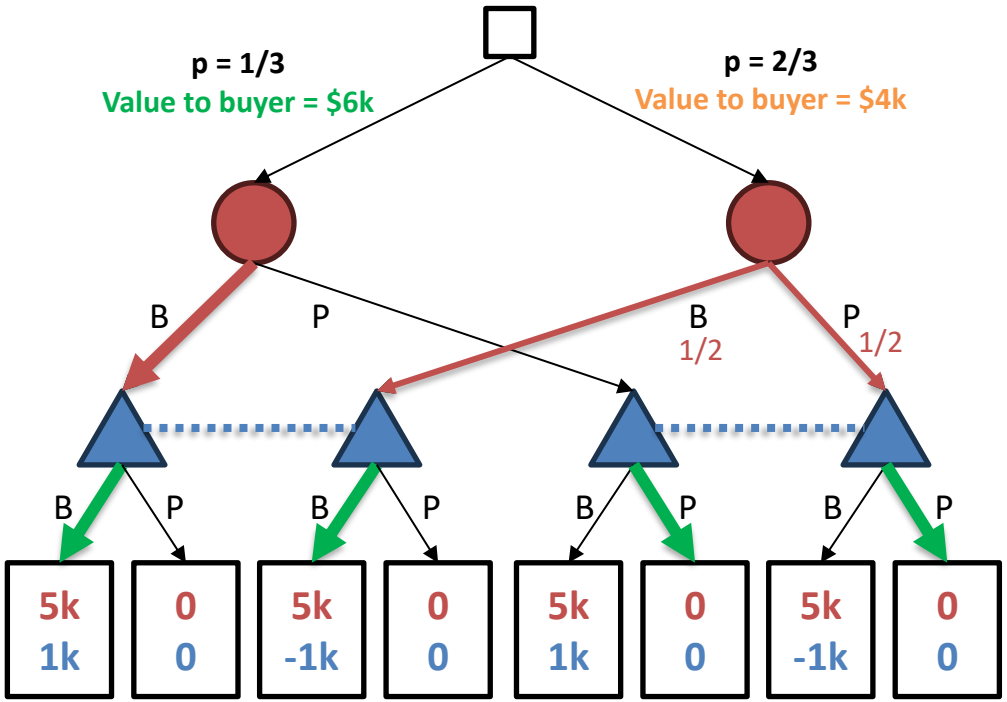
Buyer: Pr[car is good | seller says "buy"] = 1/2

⇒ 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



You should buy the car

I should buy the car



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

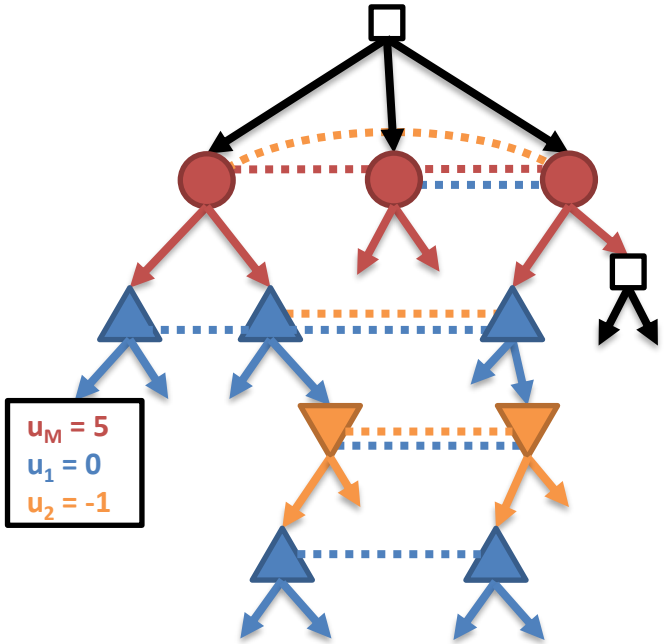
What's common to all these problems?

- **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



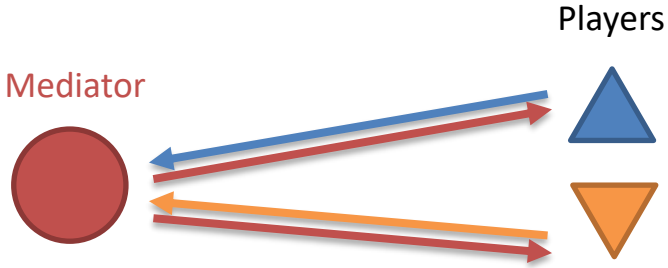
Definition:

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

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

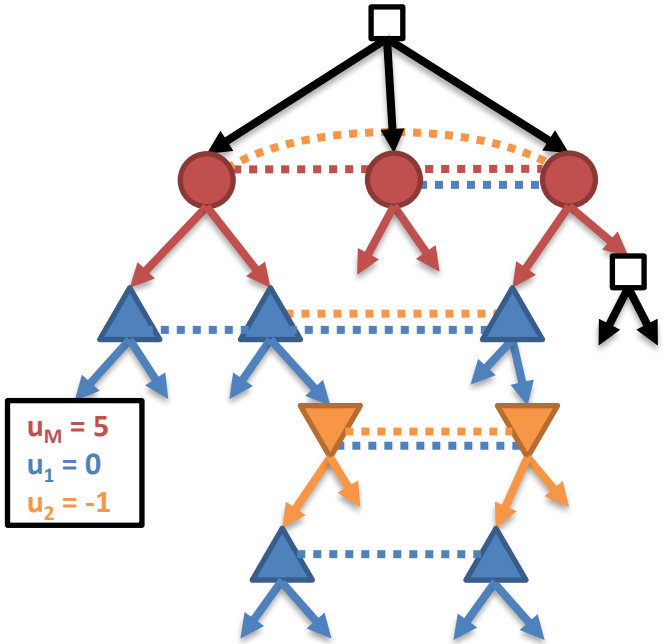
$$u_2(\mu, x_1, x_2) = \max_{x_2'} u_1(\mu, x_1, x_2')$$

Communication is implicit. At every timestep in the game…

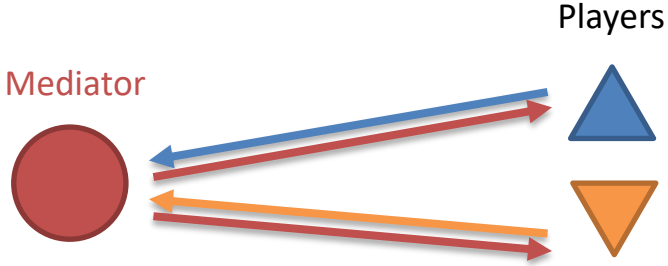


Equivalently: (x_1, x_2) is a Nash equilibrium with μ held fixed

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:

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

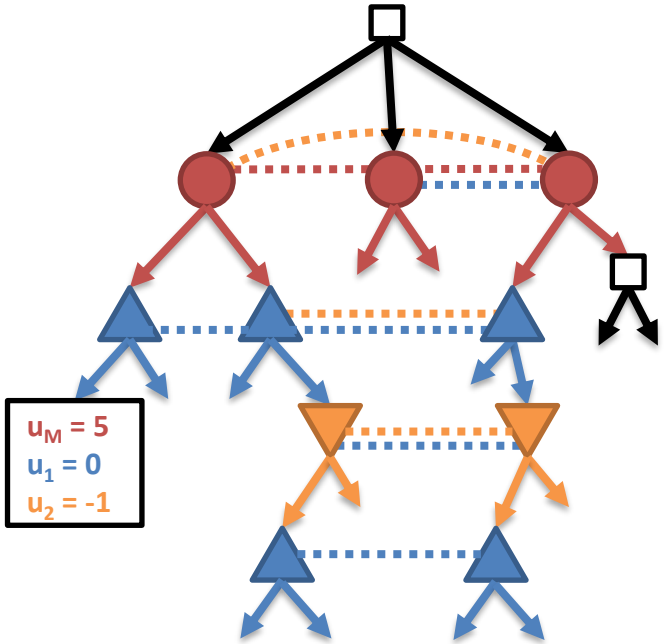
$$u_2(\mu, x_1, x_2) = \max_{x_2'} u_1(\mu, x_1, x_2')$$

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:

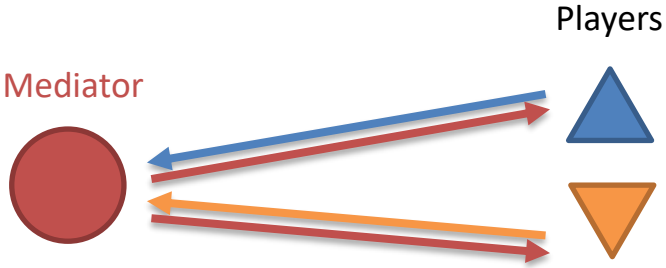
Communication equilibrium: *tuple of (possibly randomized) strategies* (μ, x_1, x_2) s.t. all **players** (not incl. mediator) are best-responding:

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

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

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

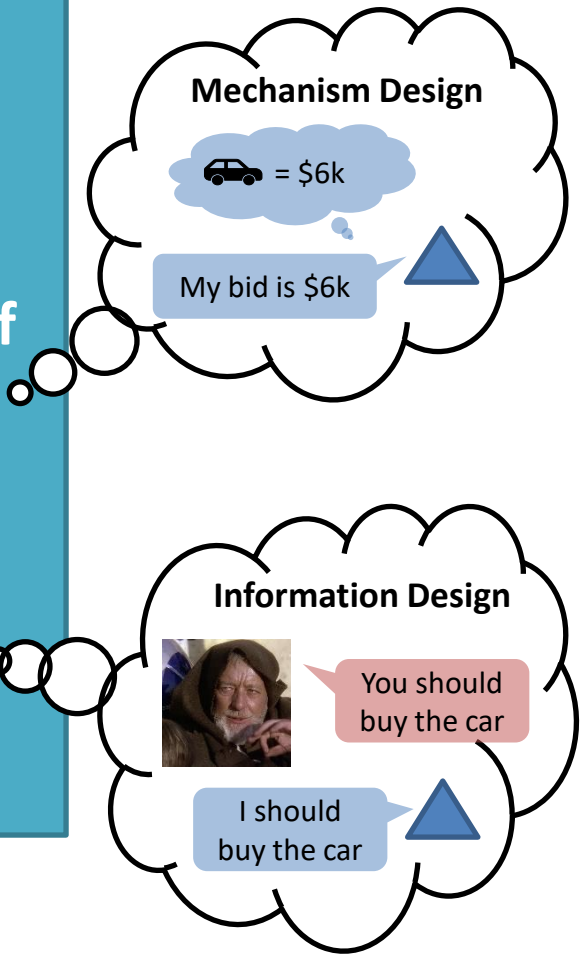
Theorem (Revelation Principle, *informal*)

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

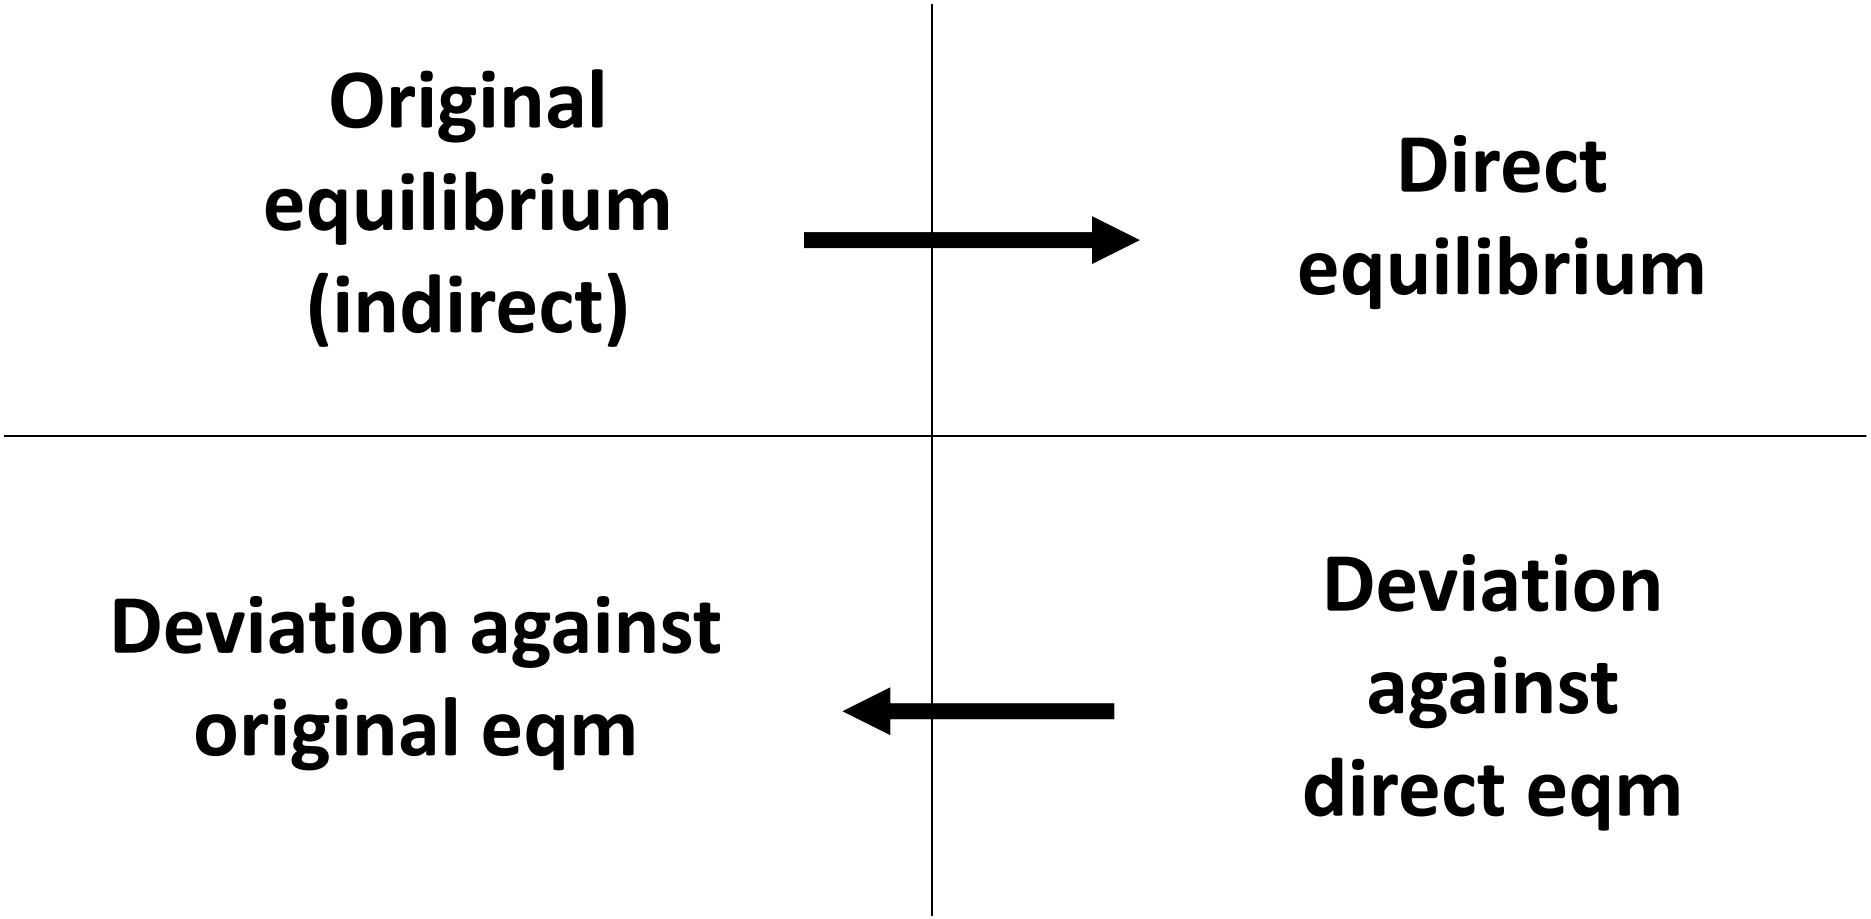
For every comm eqm, exists equivalent **direct** comm eqm. "Direct" means both:

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

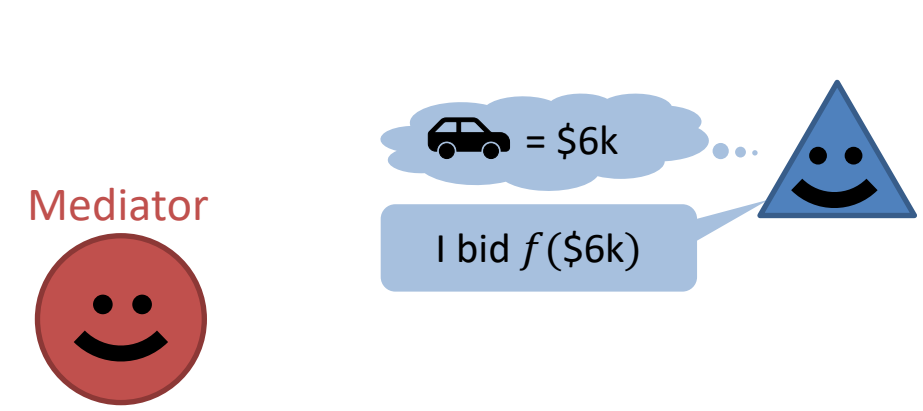


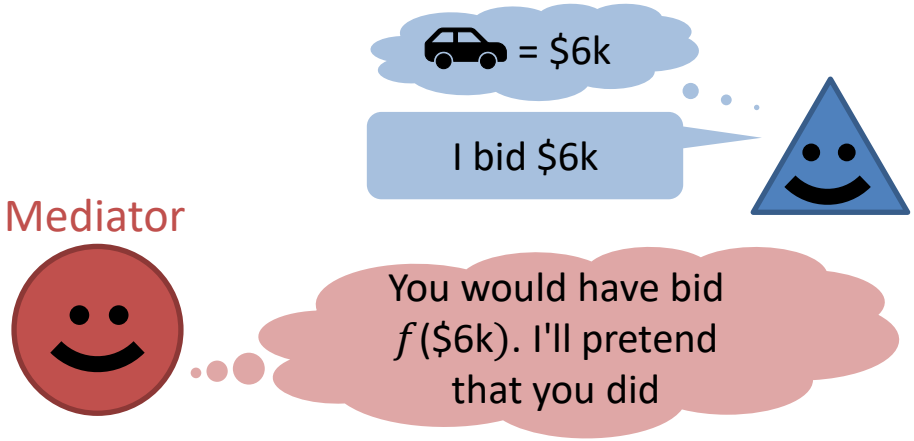
Step 1: Revelation Principle

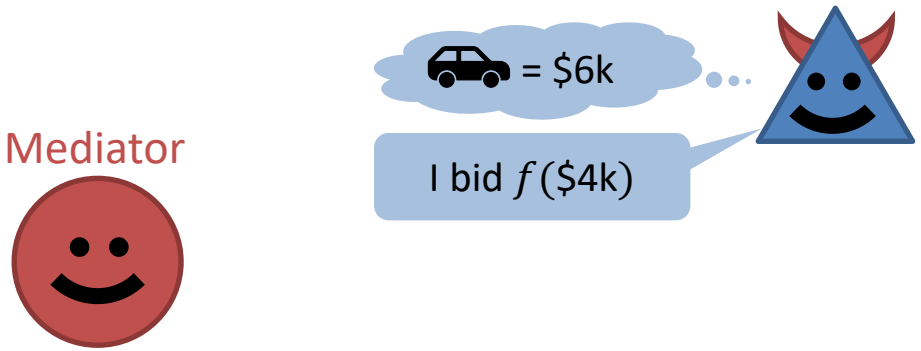
1. The players' messages to the mediator are reports of private information. In equilibrium, players always send their true private information.

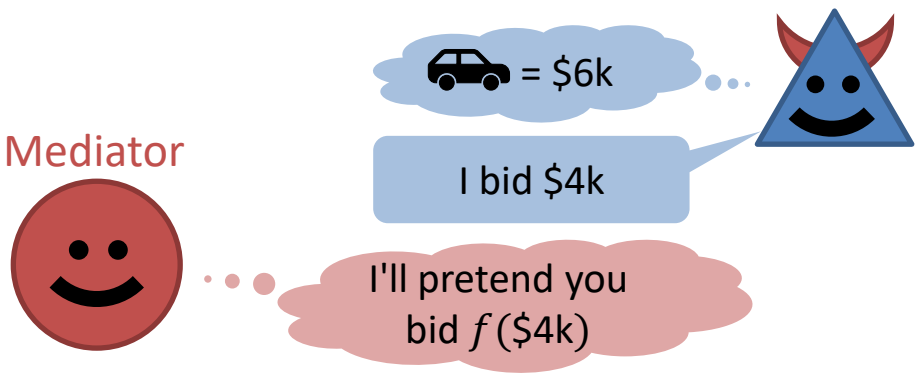
Original equilibrium (indirect)

Direct equilibrium





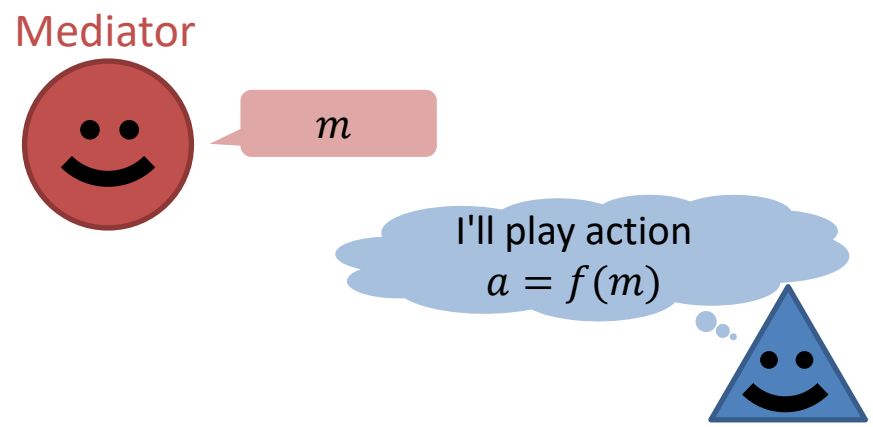




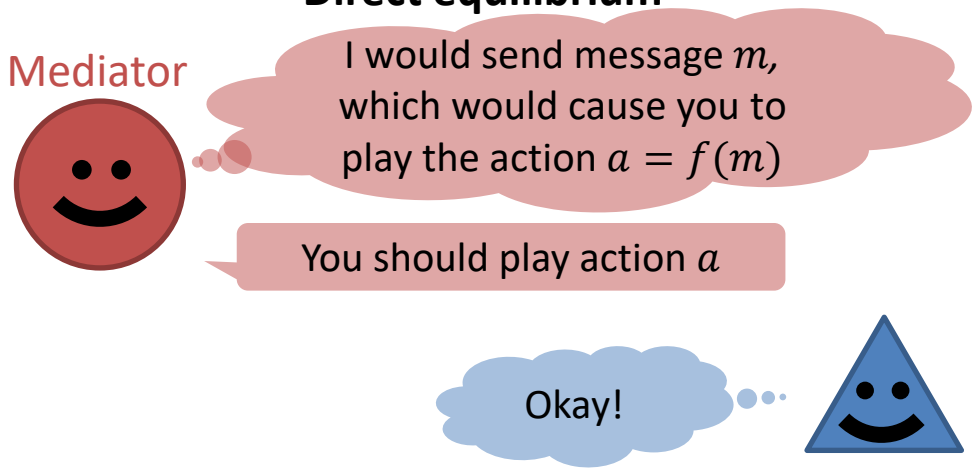
Step 1: Revelation Principle

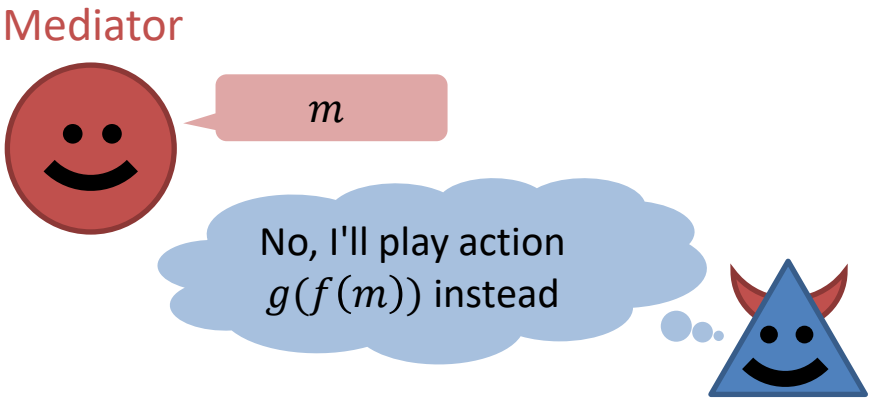
2. The mediator's messages to the players are action recommendations. In equilibrium, players play the actions that they are recommended.

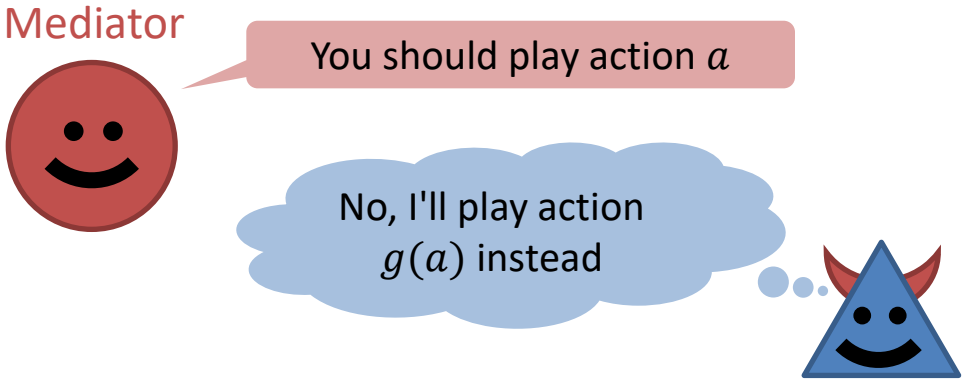
Original equilibrium (indirect)



Direct equilibrium







Proof in Three Steps

- Step 1: Reduce game from infinite to finite.

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



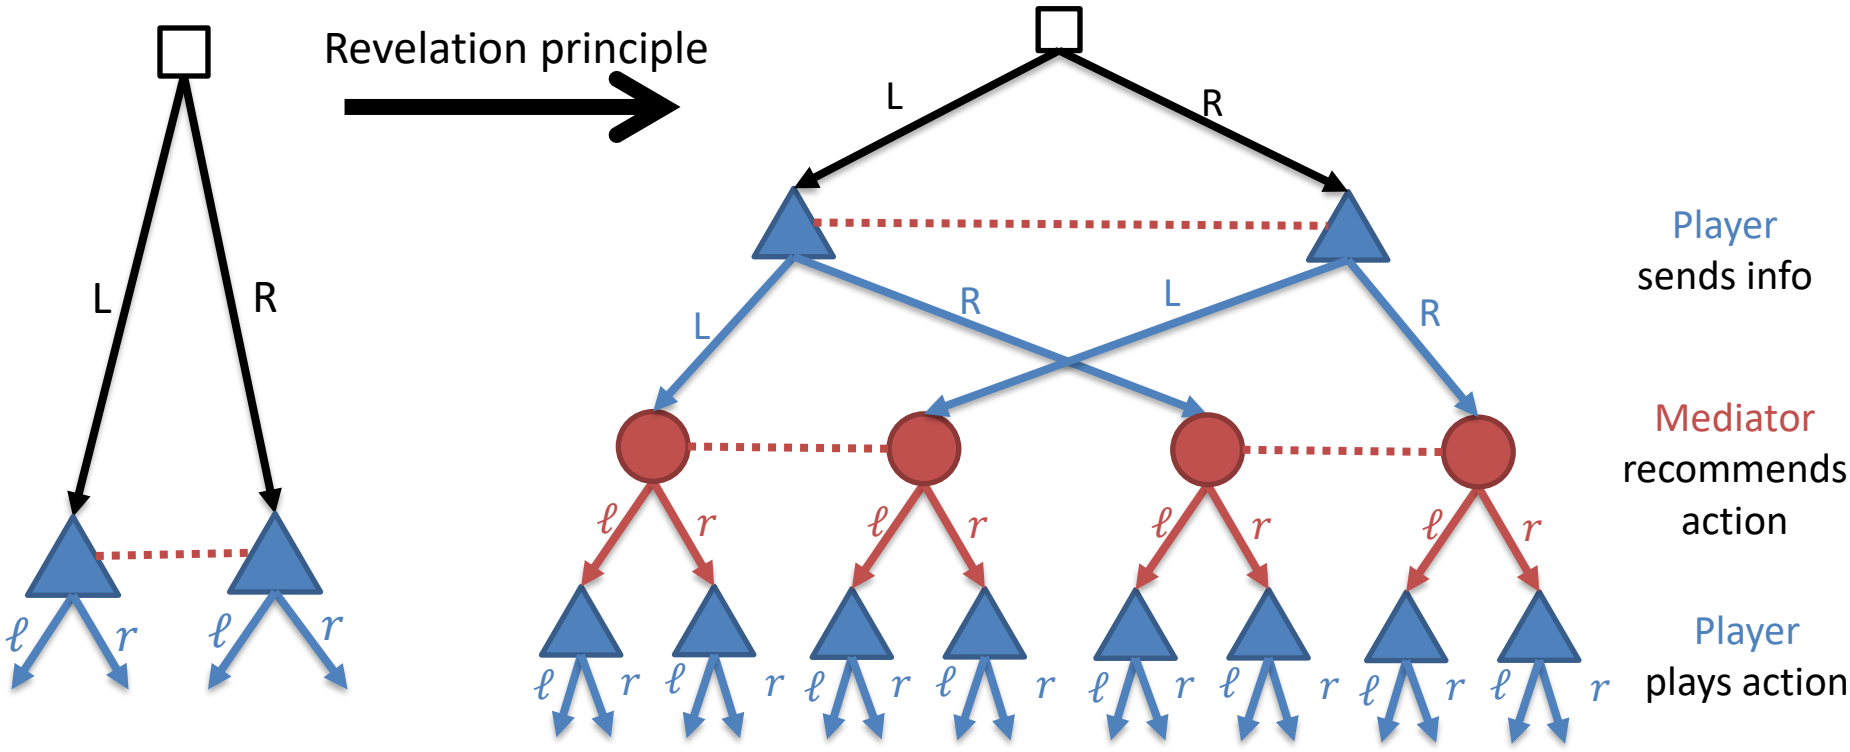
- Step 2: Reduce game from finite to polynomial.

  (Using more "without loss of generality" reductions)

- Step 3: Solve game.

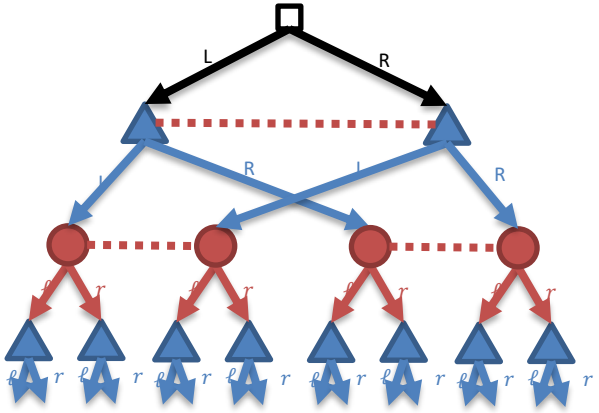
  (LP duality)

Step 2: Communication Game



"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

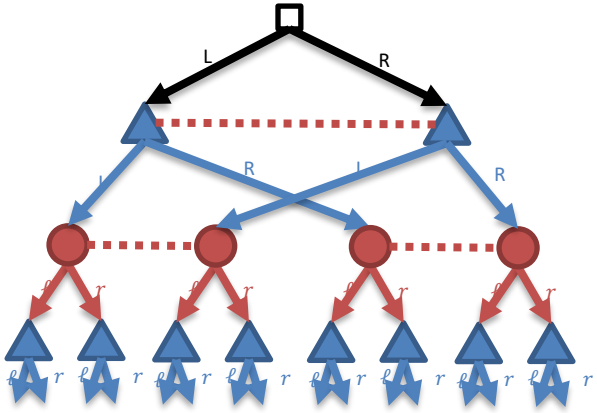
#states in communication game tree
$\leq |S| \cdot (\#\text{possible messages})^{O(\text{game tree depth}) \cdot n}$

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.
⇒ Mediator shouldn't give Player i any useful info
⇒ Players can't benefit from such τ_i

Step 2: Communication Game



State in communication game tree

$\updownarrow$

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

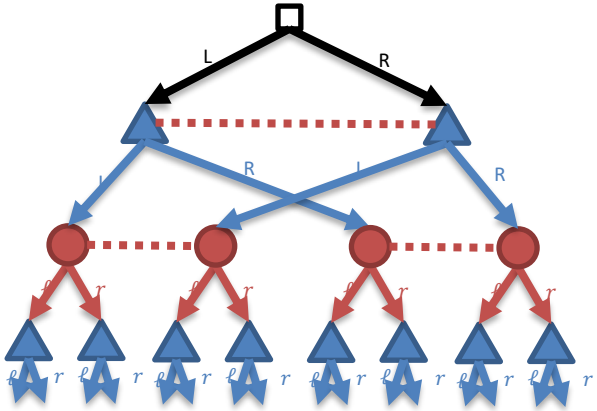
$s \in S$: state in original game tree
$s_i \in S$: state corresponding to τ_i
$n =$ # players

#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
$i \in [n] \cup \{\perp\}$: 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* $(\boldsymbol{\mu}, \boldsymbol{x_1}, \dots, \boldsymbol{x_n})$ such that all **players** (not incl. mediator) are best-responding:

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

↓ Steps 1 & 2

Definition:

A **direct communication equilibrium** 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

$$\boldsymbol{x_i^*} = \textbf{direct strategy of player } \boldsymbol{i}$$

(Send honest info, obey recommendations)

Step 3: Solving the Game

Program:

$$\max u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x_1^*}, \ldots, \boldsymbol{x_i^*}, \ldots, \boldsymbol{x_n^*})$$

linear in $\boldsymbol{\mu}$

constants

s.t.

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

linear in $\boldsymbol{\mu}$

bilinear in $\boldsymbol{\mu}, \boldsymbol{x_i'}$

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:

$$\max \; \underbrace{c^\top \mu}_{\text{linear in } \mu} \quad \text{s.t.} \quad \underbrace{b_i^\top \mu}_{\text{linear in } \mu} \geq \max_{x_i'} \underbrace{\mu^\top A_i x_i'}_{\text{bilinear in } \mu, x_i'} \text{ for all } i$$

Definition:

A **(direct) comm eq** 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

$$x_i^* = \textbf{direct strategy of player } i$$

(Send honest info, obey recommendations)

Step 3: Solving the Game

Program:

$$\max_{\mu \in X_M} c^\top \mu \quad \text{s.t.}$$

$$b_i^\top \mu \geq \max_{x_i' \in X_i} \mu^\top A_i x_i' \text{ for all } i$$

Take duals of inner maximizations

Let $x_i' \in X_i = \{x : F_i x = f_i, x \geq 0\}$



Linear program:

$$\max_{\mu \in X_M,\, v_i : i \in [n]} c^\top \mu \quad \text{s.t.}$$

$$b_i^\top \mu \geq f_i^\top v_i, \quad F_i^\top v_i \geq A_i^\top \mu \quad \text{for all } i$$



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:

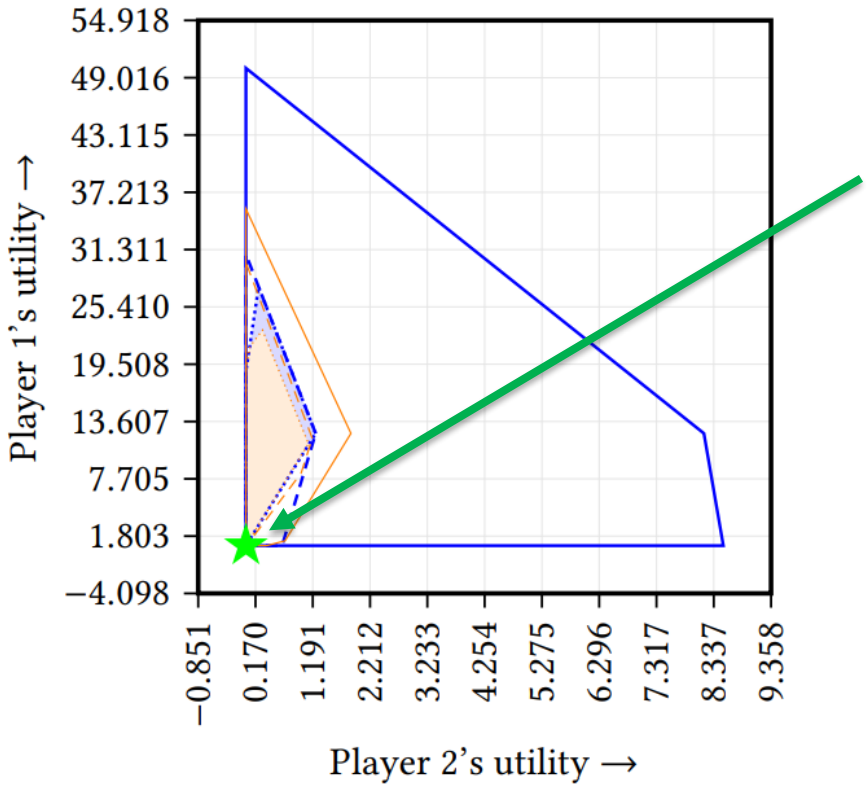
- Optimal sequential mechanism design
- Optimal sequential information design

…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!



Other notions of equilibrium

★ Communication equilibrium

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

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

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

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

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

This is a zero-sum game!

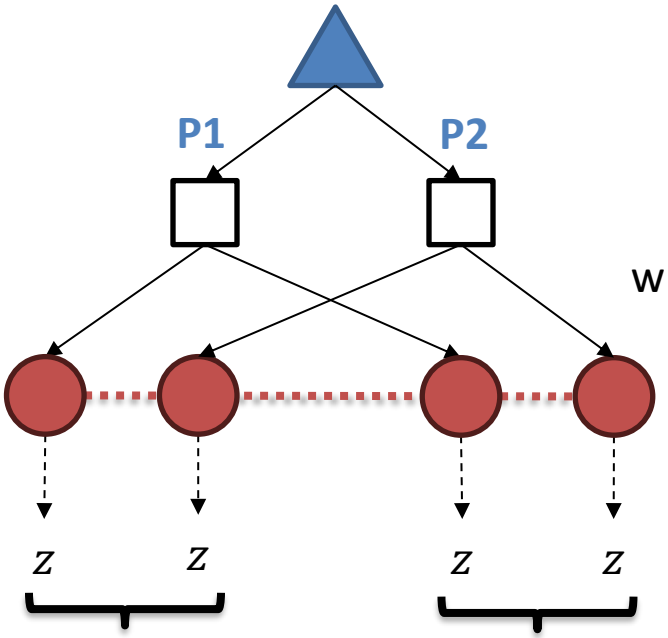
Proposition: There exists $\lambda^* > 0$ s.t. for all $\lambda > \lambda^*$:

Equilibrium strategy for max-player of this zero-sum game $\equiv$ Optimal correlated 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 correlated 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 $\lambda := 1/\varepsilon$

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

$$\max_{\boldsymbol{\mu}} u_{\mathrm{M}}(\boldsymbol{\mu}, \boldsymbol{x}^*) \quad \text{s.t.}$$

for all players i

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



find $\boldsymbol{\mu}$ s.t.

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

and for all players i

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

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 $\boldsymbol{\mu}$ guaranteeing value $> -\varepsilon$ (branch high), or
- it proves value < 0 (branch low)

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

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

Solution #2: An Alternative Lagrangian

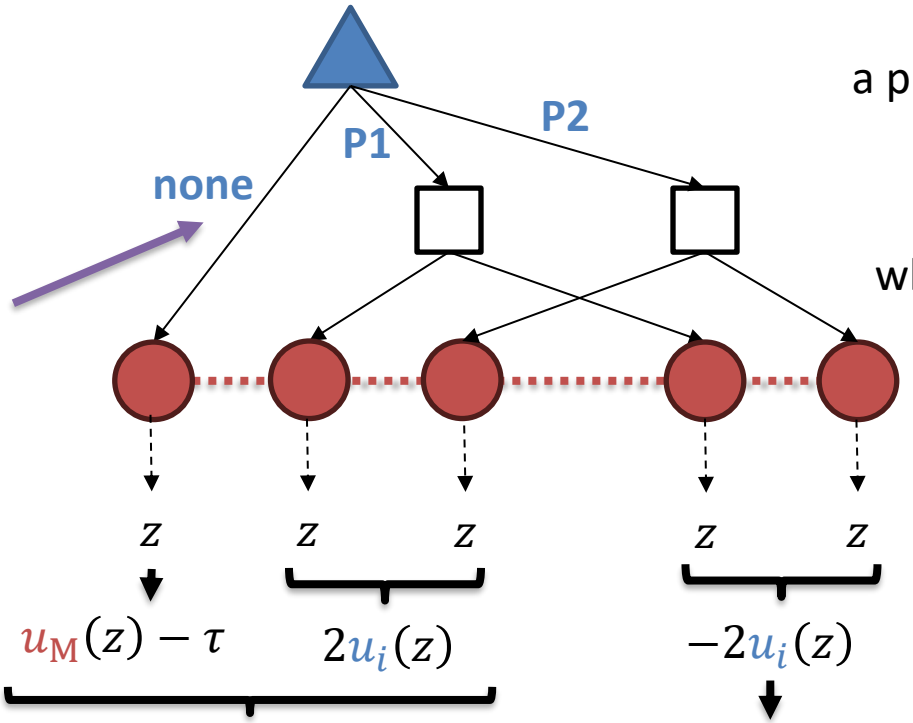
Theorem:
The last $\boldsymbol{\mu}$ found by the binary search algorithm is an ε-equilibrium whose mediator objective is at least $v^* - O(\varepsilon)$
(where v^* = optimal equilibrium mediator objective)

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 $\geq \tau$

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

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:

Mediator plays with all **direct players**

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

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

Which is Better?

"Direct" Lagrangian

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

☺ 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

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

☹ 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



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

Amenable to deep RL!

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