Team games

Brian Zhang

What are team games?

P1 and P2 are **teammates**
(same utility function)

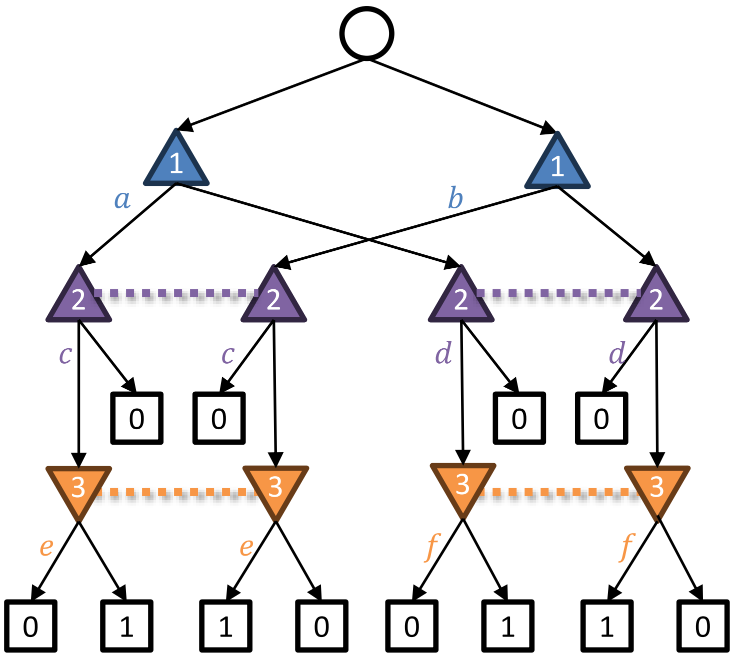

"Communication game": P1 must send one bit of information to P2 without P3 knowing

Goal: Find an optimal **strategy profile for the team**

$$\max_{\substack{x_1\in\Delta(\Pi_1)\\ x_2\in\Delta(\Pi_2)}} \min_{y\in\Delta(\Pi_3)} u(x_1, x_2, y)$$

$$\max_{\substack{x_1\in\Delta(\Pi_1)\\ x_2\in\Delta(\Pi_2)}} \min_{x_3\in\Delta(\Pi_3)} u(x_1, x_2, y)$$

$\bar{x} \coloneqq 1 - x$

$$= \max_{a,b,c,d\in[0,1]} \min_{e,f\in[0,1]} \frac{1}{2}\left(ac\bar{e} + \bar{a}d\bar{f} + b\bar{c}e + \bar{b}\bar{d}f\right)$$

$$\le \max_{a,b,c,d} \frac{1}{2}\left(ac\bar{c} + \bar{a}d\bar{d} + b\bar{c}c + \bar{b}\bar{d}d\right)$$

$e \leftarrow c,$
$f \leftarrow d$

$$\le \frac{1}{2}\cdot\frac{1}{4}\cdot\left(a + \bar{a} + b + \bar{b}\right) = \frac{1}{4}$$

$x\bar{x} = x(1-x) \le \frac{1}{4}$

(in fact, $= 1/4$ by setting $a = b = c = d = 1/2$)

What are team games?

P1 and P2 are **teammates**
(same utility function)

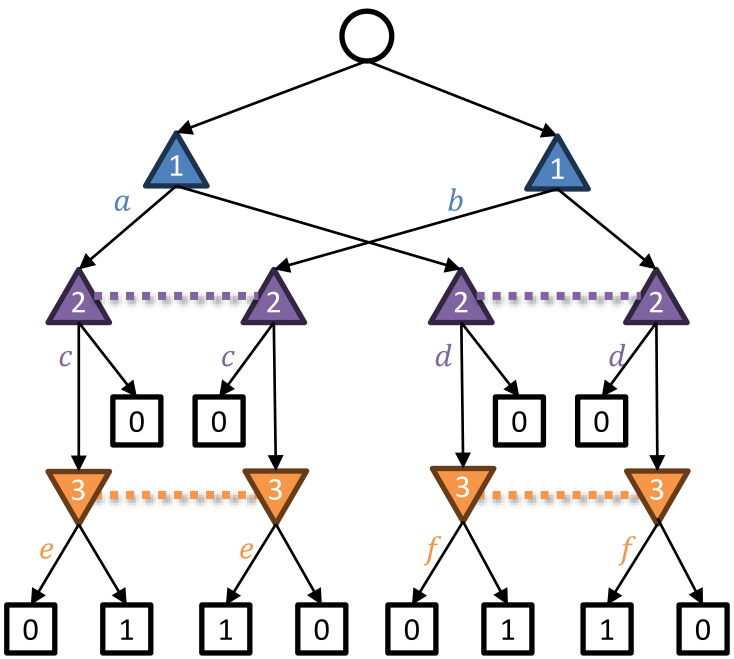

$$\min_{y\in\Delta(\Pi_3)} \max_{\substack{x_1\in\Delta(\Pi_1)\\ x_2\in\Delta(\Pi_2)}} u(x_1, x_2, y)$$

$$= \min_{e,f} \max_{a,b,c,d} \frac{1}{2}\left(ac\bar{e} + \bar{a}d\bar{f} + b\bar{c}e + \bar{b}\bar{d}f\right)$$

$$\geq \frac{1}{2} \min_{e,f} \max\{\bar{f} + e, \bar{e} + f\}$$

either $a = c = 0,\ b = d = 1$
or $a = c = 1,\ b = d = 0$

$$= \frac{1}{2}$$

set $e = f$ so that $\bar{f} + e = \bar{e} + f = 1$

$$> \frac{1}{4}$$

Why?

"Communication game": P1 must send one bit of information to P2 without P3 knowing

Goal: Find an optimal **strategy profile for the team**

$$\min_{y\in\Delta(\Pi_3)} \max_{\substack{x_1\in\Delta(\Pi_1)\\ x_2\in\Delta(\Pi_2)}} u(x_1, x_2, y) \quad ?$$

Solution concepts for team games

P1 and P2 are **teammates**
(same utility function)

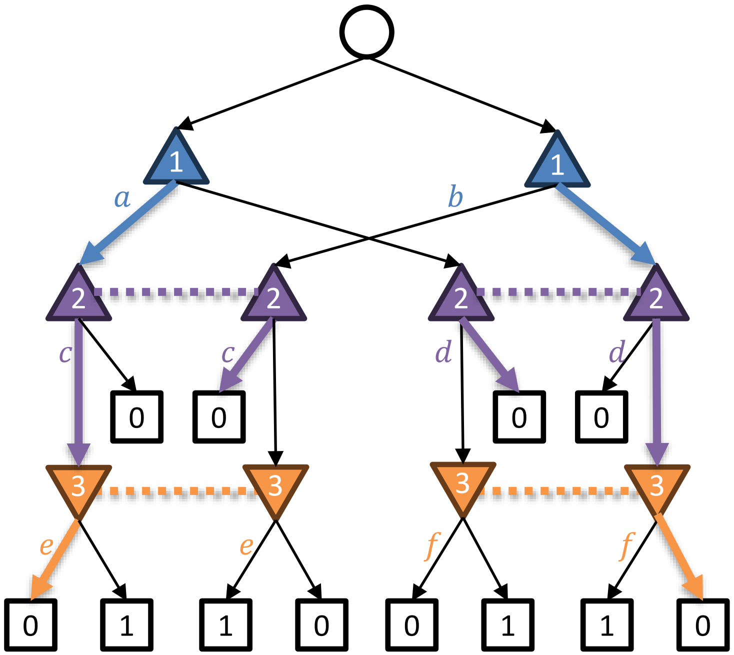

"Communication game": P1 must send one bit of information to P2 without P3 knowing

Nash equilibrium?

Bad: $(a, b; c, d; e, f) = (1,0; 1,0; 1,0)$ is a Nash eq. in which P1/P2 always lose

Solution concepts for team games

P1 and P2 are **teammates**
(same utility function)

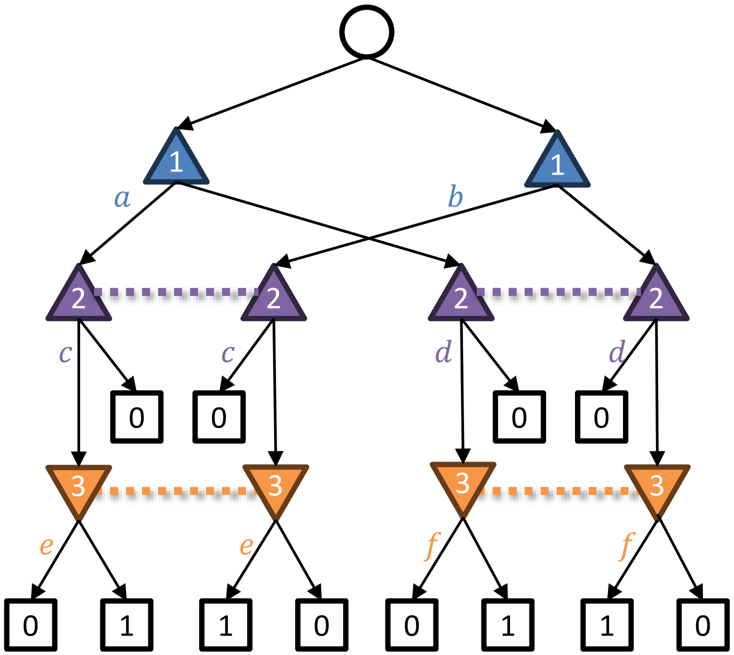

"Communication game": P1 must send one bit of information to P2 without P3 knowing

Nash equilibrium?

Bad: $(a, b; c, d; e, f) = (1,0; 1,0; 1,0)$ is a Nash eq. in which P1/P2 always lose

Team min-max equilibrium (TME):

Solutions to

$$\max_{\substack{x_1 \in \Delta(\Pi_1) \\ x_2 \in \Delta(\Pi_2)}} \min_{y \in \Delta(\Pi_3)} u(x_1, x_2, y)$$

(These are what we computed.)

Problem: Nonconvex optimization
⇒ hard; minimax theorem doesn't apply

"Equilibrium"?

Solution concepts for team games

P1 and P2 are **teammates**
(same utility function)

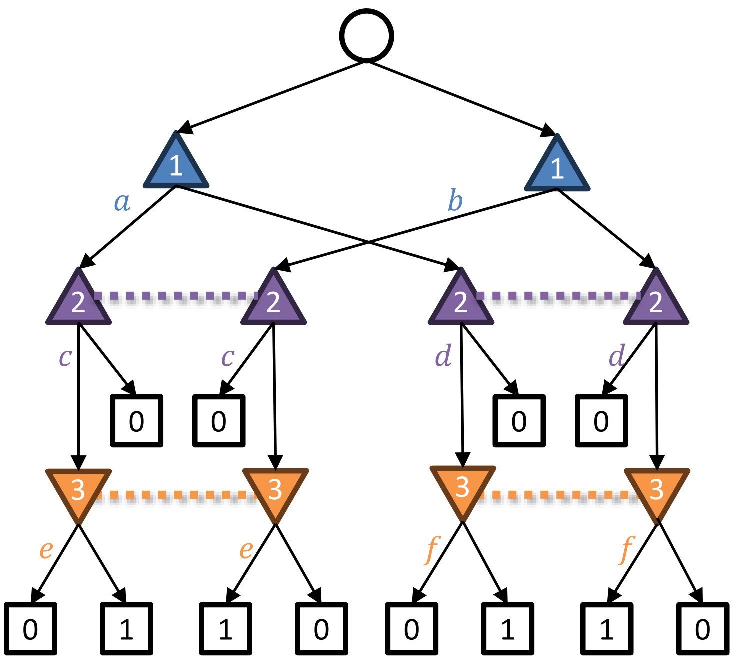

"Communication game": P1 must send one bit of information to P2 without P3 knowing

Team min-max equilibrium **with correlation (TMECor)**:

Solutions to

$$\max_{\mu\in\Delta(\Pi_1\times\Pi_2)} \min_{y\in\Delta(\Pi_3)} \mathbb{E}_{(x_1,x_2)\sim\mu} u(x_1,x_2,y)$$

Players on the same team can **correlate** their strategies using **team-private randomness**

In this game:

- TME value = 1/4
- TMECor value = 1/2: P1/P2 randomize uniformly between (1,0; 1,0) and (0,1; 0,1) ("random one-time pad")

TME	TMECor
nonconvex optimization problem	convex optimization problem (bilinear saddle point)
minimax theorem fails	minimax theorem holds

Equivalently: (Timeable) two-player zero-sum games of imperfect recall

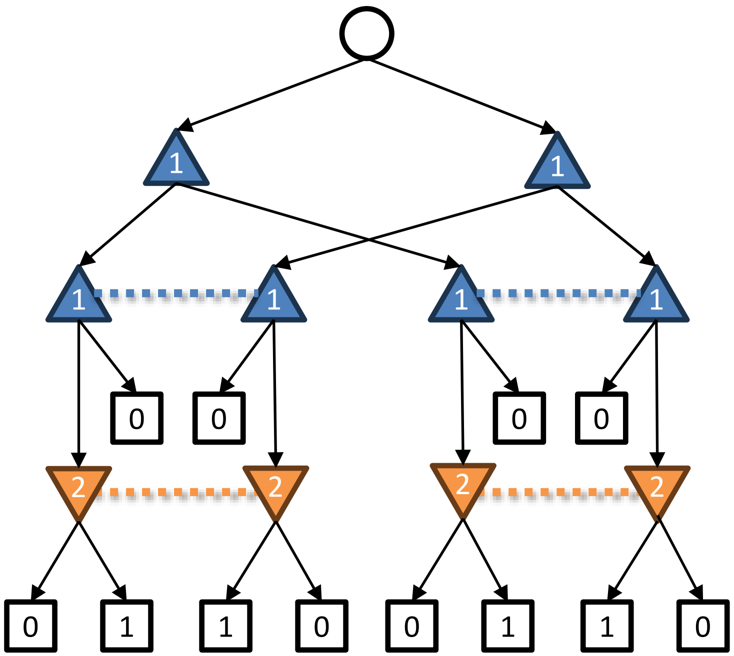

Team games	Imperfect-recall games
team	player
strategy profile for team	strategy for player
uncorrelated strategy profile	behavioral strategy
correlated strategy profile	mixed strategy
TME	optimal behavioral strategy
TMECor	mixed Nash equilibrium

⚠ Kuhn's theorem does not apply!
(mixed ≠ behavioral for imperfect-recall games)

Question for remainder of class: When can we **compute TMECor efficiently?**

Representing strategy profiles of a team: Realization form

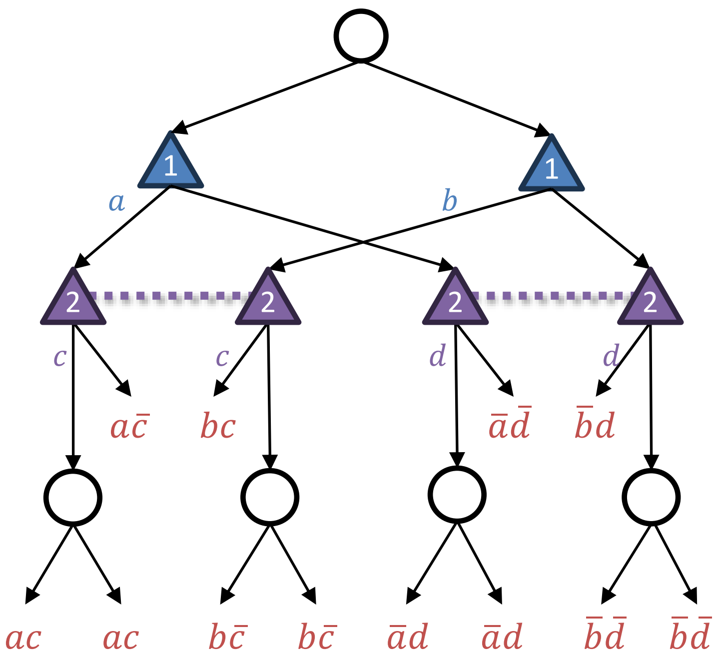

Pure strategies:
$\boldsymbol{x} \in \{0,1\}^Z$ ($Z =$ set of terminal nodes)
$x[z] = 1$ if **team ▲** plays all actions on path to z
(similar for team ▼)

TMECor: saddle-point solutions to

$$\max_{x \in \operatorname{conv} X} \min_{y \in \operatorname{conv} Y} \boldsymbol{x}^\top \boldsymbol{A} \boldsymbol{y}$$

where

$$A[z,z] = p_{\text{chance}}(z) \cdot u(z)$$

This is a **convex, bilinear saddle-point problem!**

Intuition: if both teams have one player, this is the same as the sequence-form max-min problem (modulo a linear transformation)

Computing TMECor is hard

$$\phi(x_1, x_2, x_3) = (x_1 \lor x_2) \land (\neg x_1 \lor x_3)$$

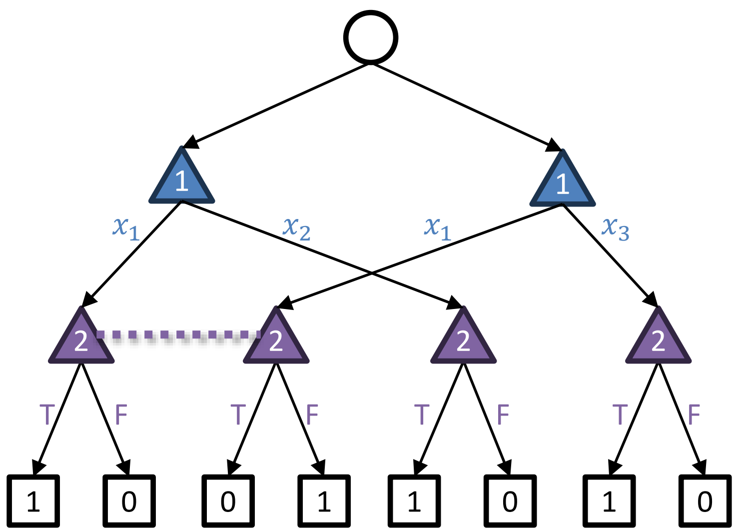

chance picks the clause

P1 picks a variable in that clause

P2 learns the variable (but not the clause) and assigns T or F

P2 pure strategies = assignments $\boldsymbol{x} \in \{\mathrm{T}, \mathrm{F}\}^3$
P1 and P2 can win w.p. 1
iff P2 plays a satisfying assignment

⇒ Even computing an optimal strategy for **one team** is NP-hard
More precisely: Computing TMECor is Δ_2^P-complete [Zhang, Farina, Sandholm ICML 2023]
⇒ Realization polytopes admit **no efficient representation**

Can we do *anything*?

Idea #1: Run PSRO.

Problem: Team best responses are hard…

But maybe for sufficiently large games, this is the only possible thing

McAleer, Farina, Zhou, Wang, Yang, Sandholm (NeurIPS 2023)

Can we do *anything*?

Idea: Construct single-player decision problem (like sequence form) by representing what is **common knowledge** for the team

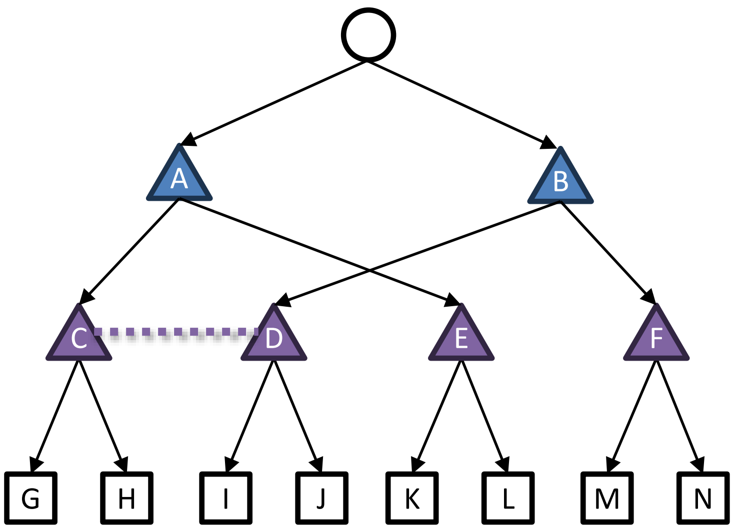

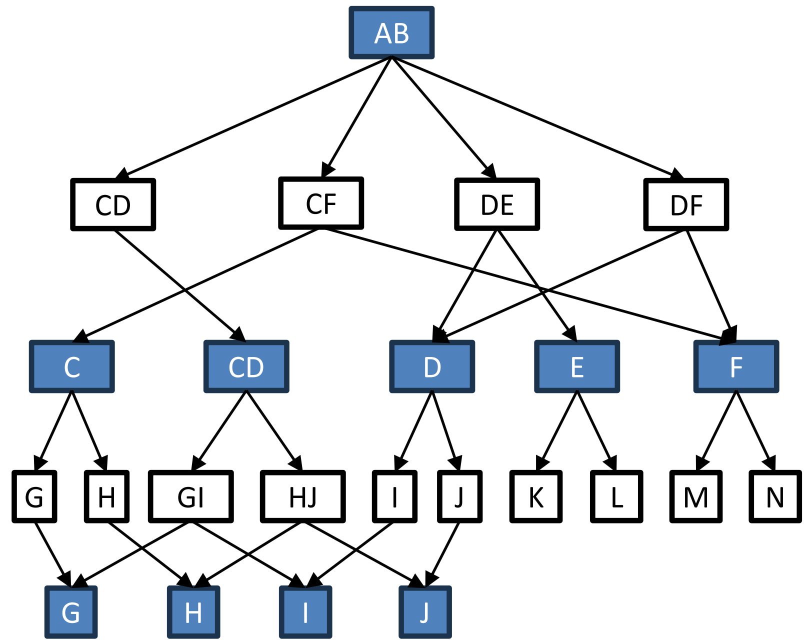

Can we do *anything*?

Idea: Construct single-player decision problem (like sequence form) by representing what is **common knowledge** for the team

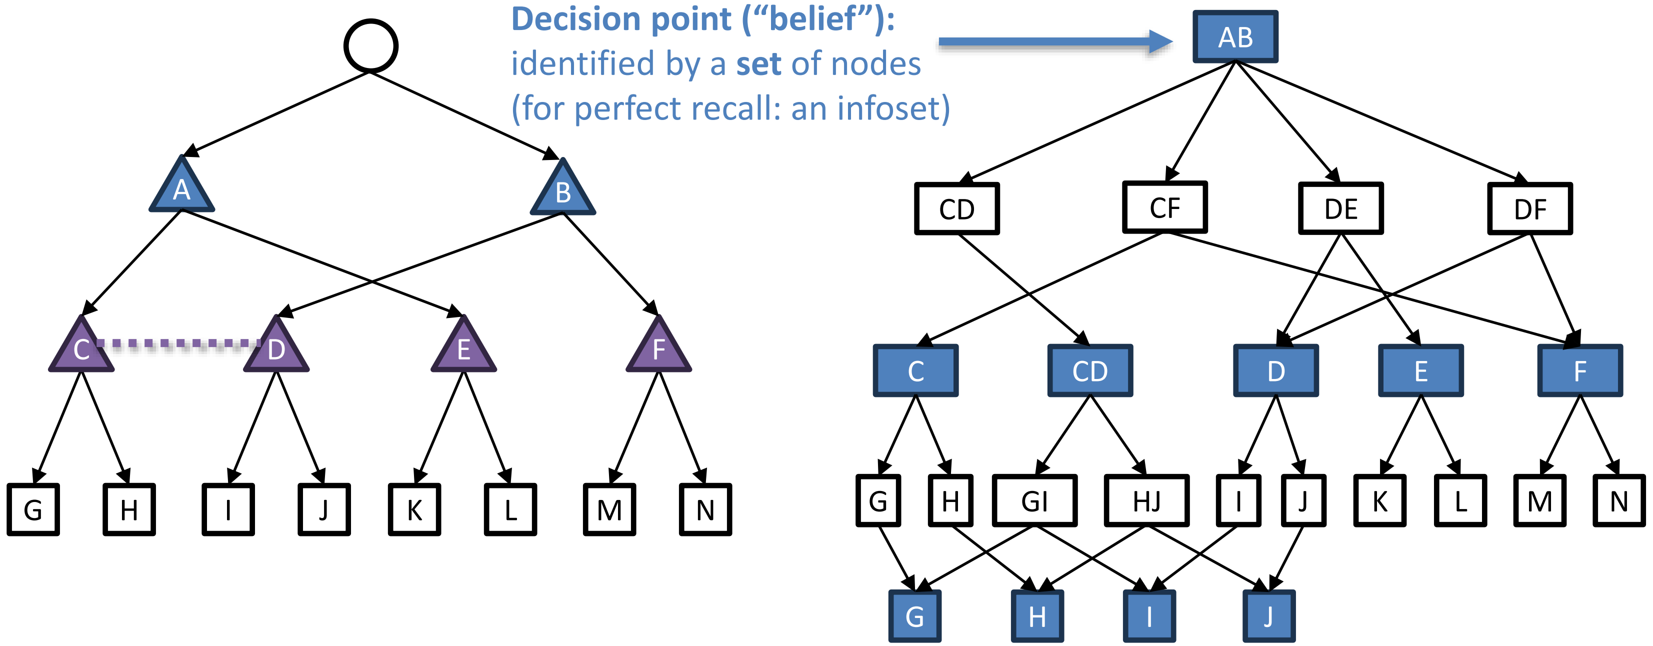

Can we do *anything*?

Idea: Construct single-player decision problem (like sequence form) by representing what is **common knowledge** for the team

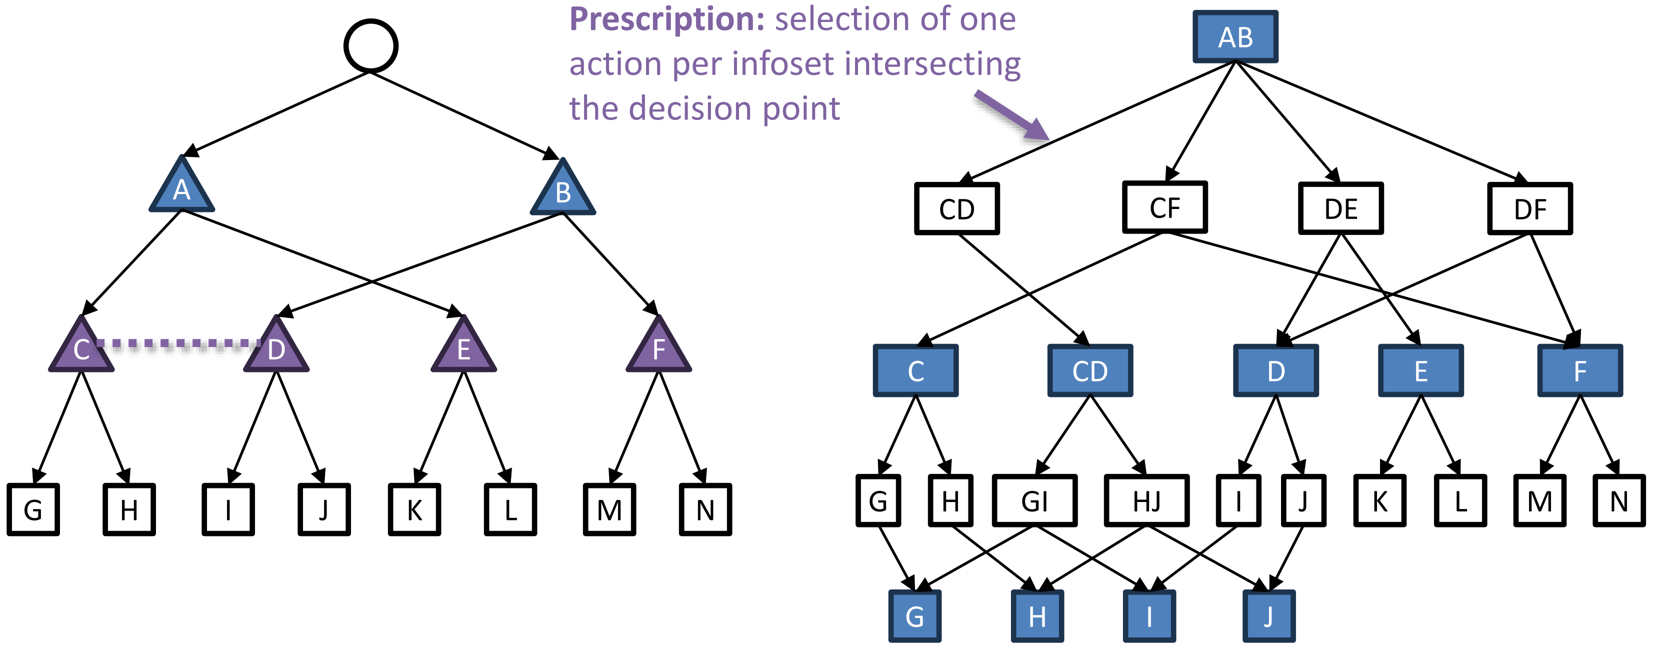

Can we do *anything*?

Idea: Construct single-player decision problem (like sequence form) by representing what is **common knowledge** for the team

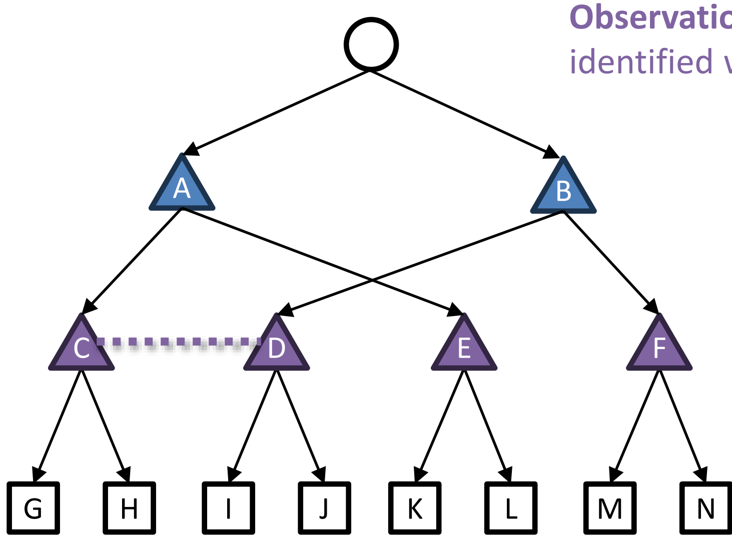

Observation point: also identified with set of nodes

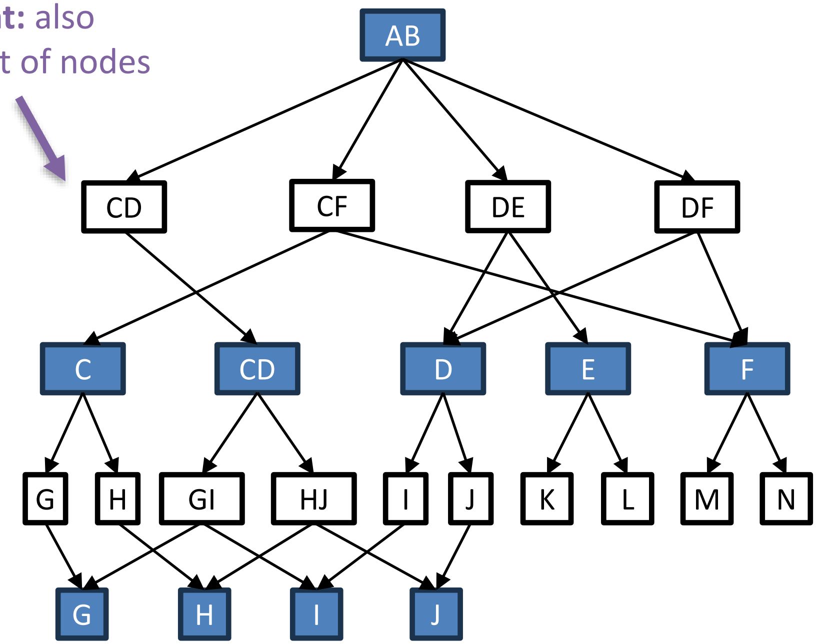

Can we do *anything*?

Idea: Construct single-player decision problem (like sequence form) by representing what is **common knowledge** for the team

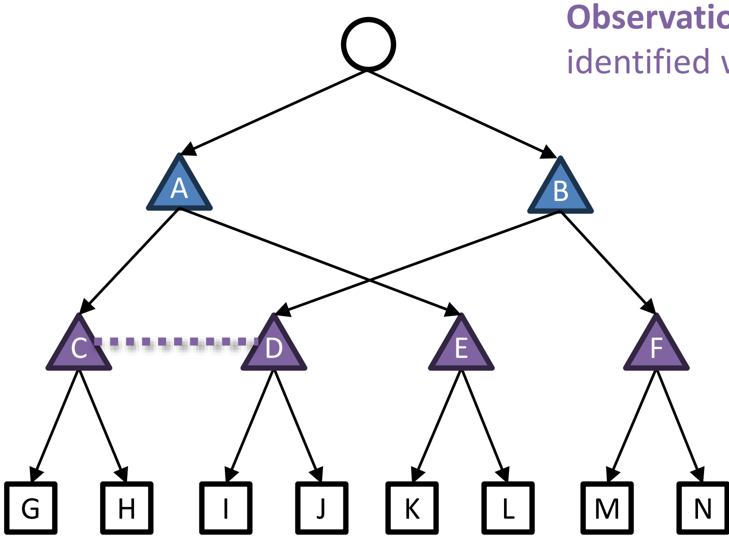

Observation point: also identified with set of nodes

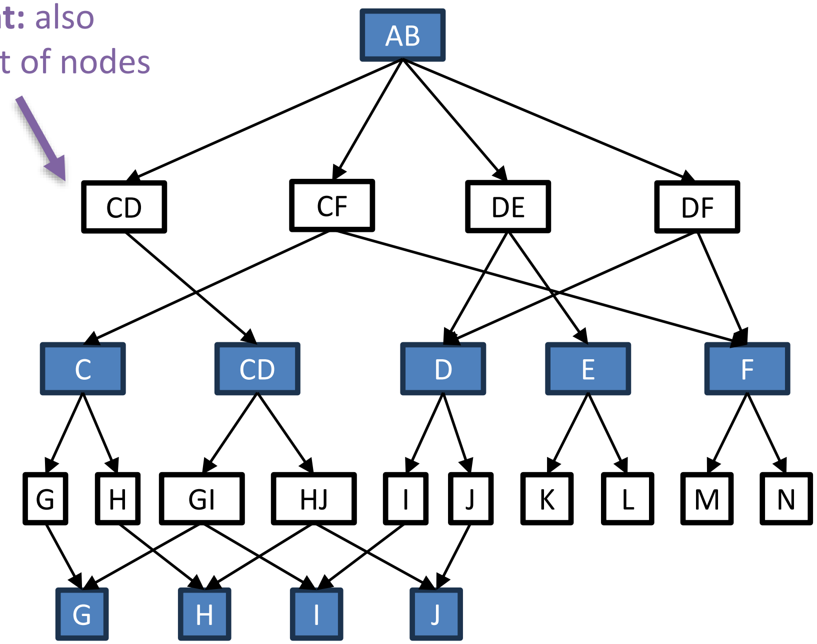

Can we do *anything*?

Idea: Construct single-player decision problem (like sequence form) by representing what is **common knowledge** for the team

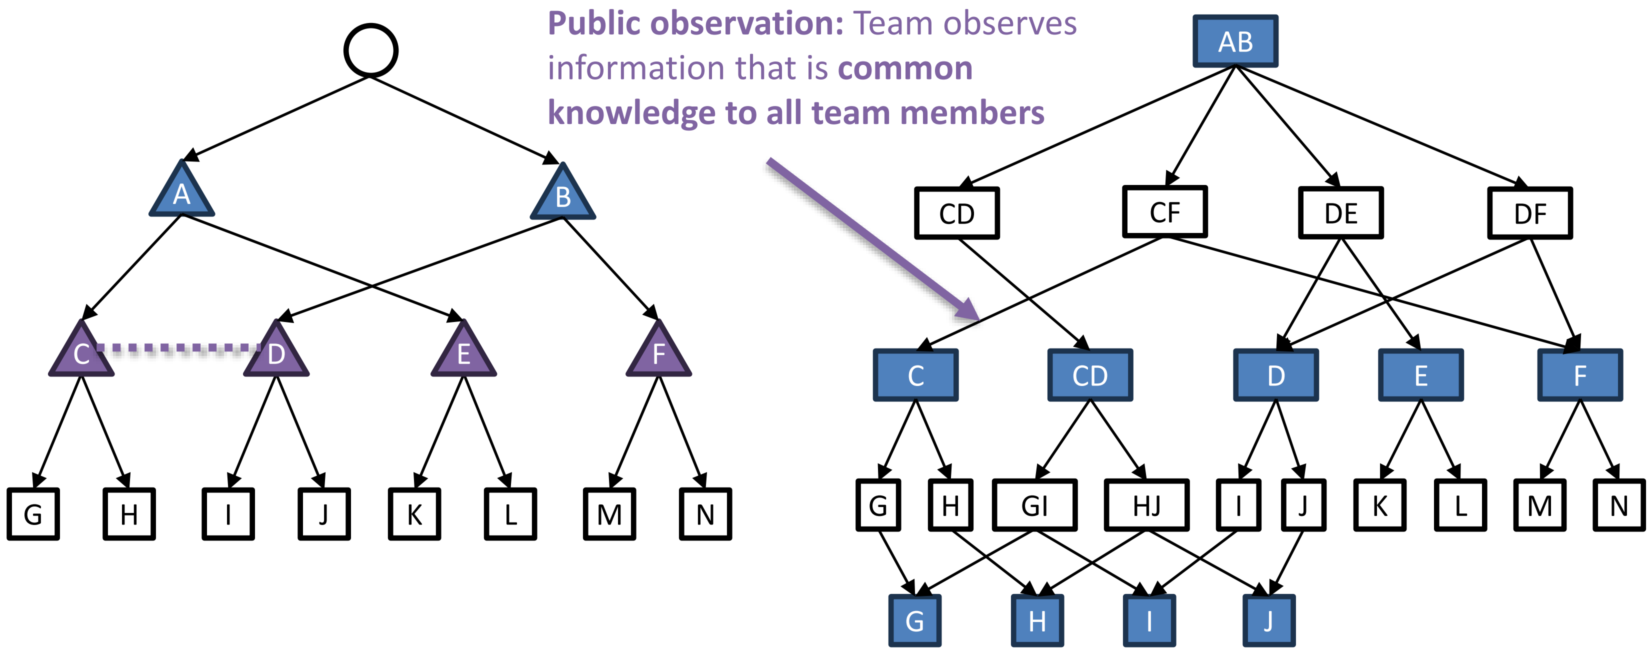

Can we do *anything*?

Idea: Construct single-player decision problem (like sequence form) by representing what is **common knowledge** for the team

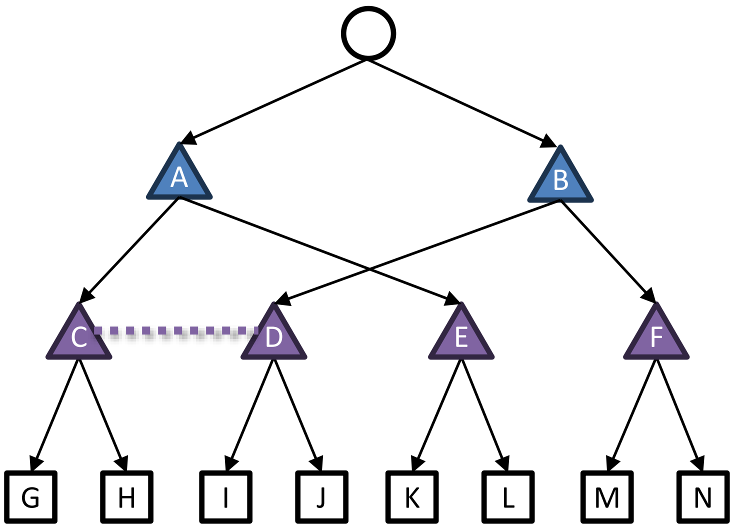

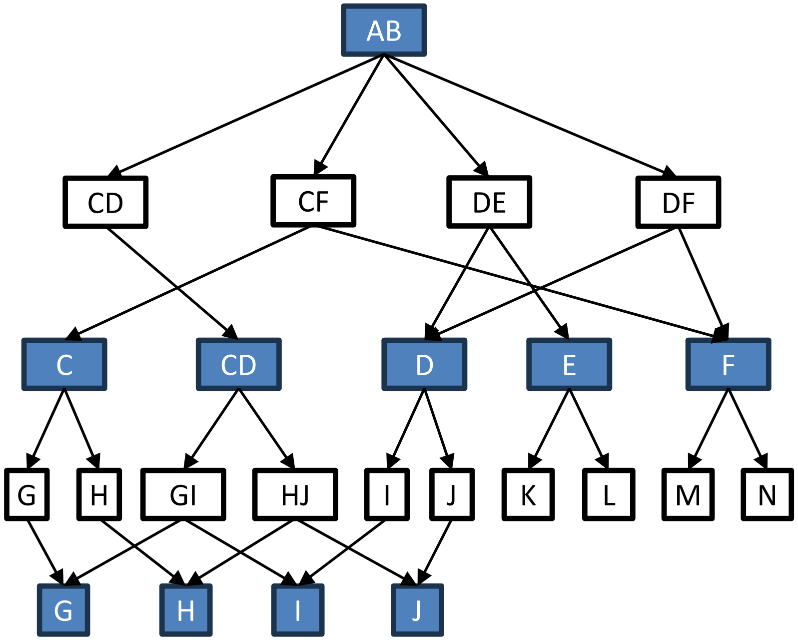

DAG decision problem: can be represented with scaled extensions (cf. linear deviations in Φ-regret, where the same idea came up!)

"Team belief DAG" (TB-DAG)

How big is the TB-DAG?

Size is dominated by observation points.

How many observation points (*i.e.,* (decision point, prescription) pairs) are there?

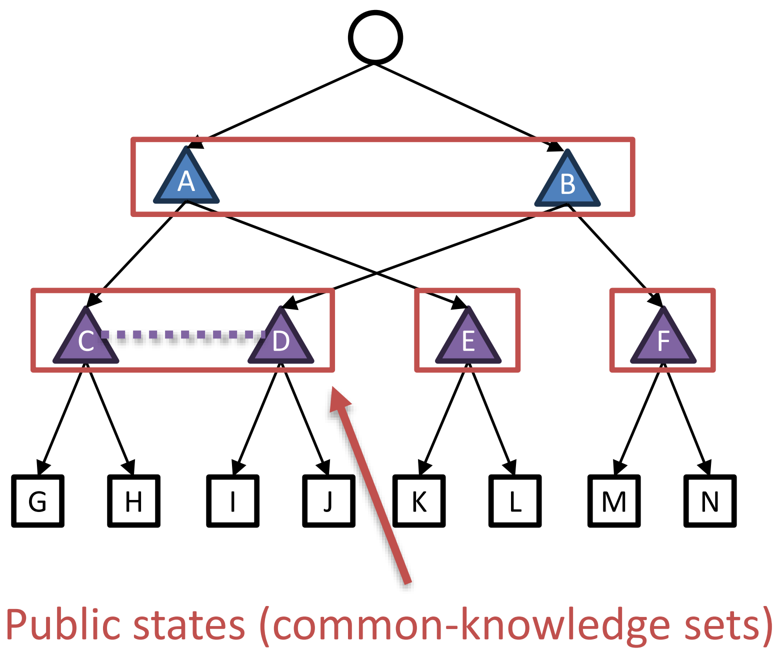

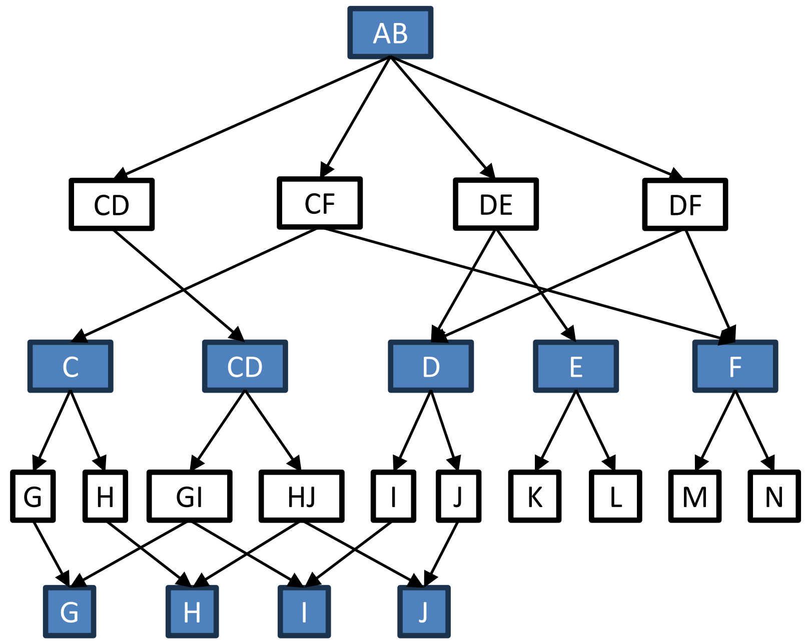

How big is the TB-DAG?

Size is dominated by observation points.
How many observation points (*i.e.*, (decision point, prescription) pairs) are there?

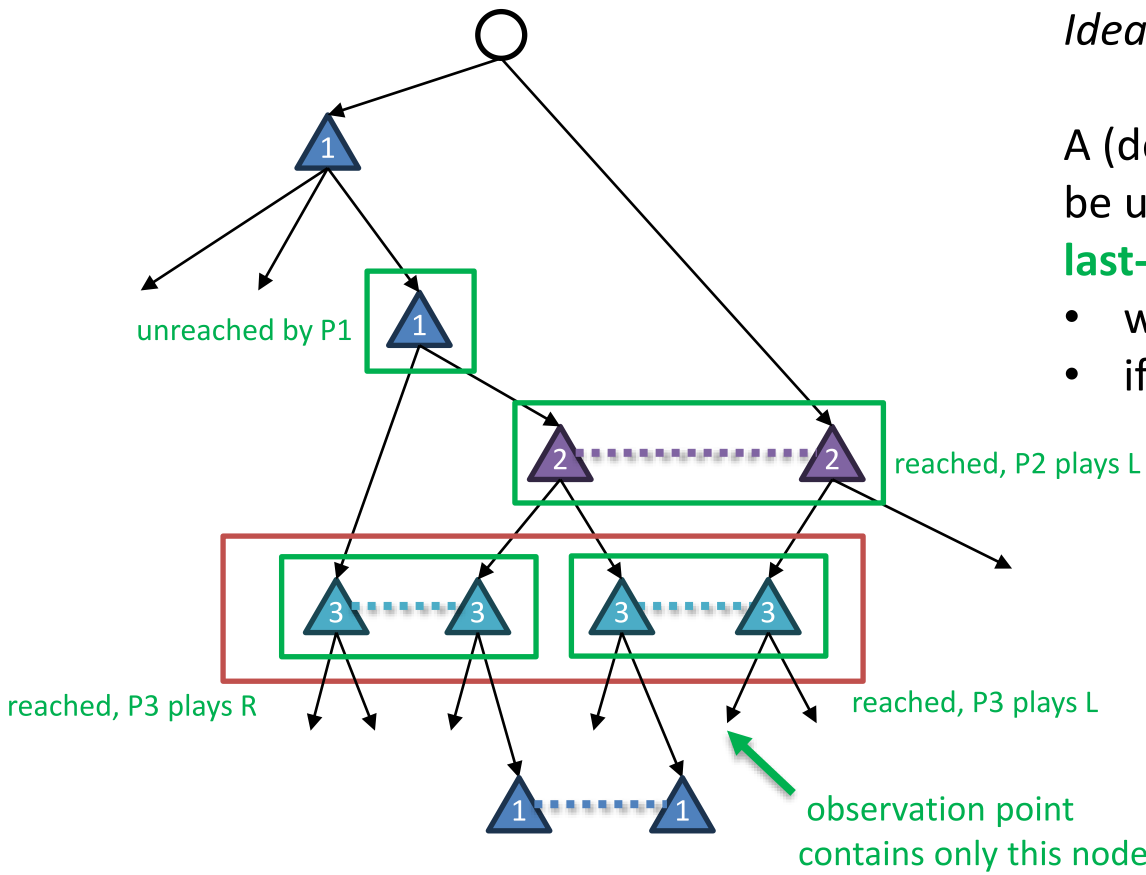

Idea: Count within each public state.

A (decision point, prescription) pair can be uniquely determined by, for each **last-infoset J,** specifying

- whether J is reached by that player
- if so, what action is selected at J

number of last-infosets

branching factor

observation points per public state $\leq (b+1)^k$

How big is the TB-DAG?

Size is dominated by observation points.

How many observation points (*i.e.*, (decision point, prescription) pairs) are there?

number of last-infosets

branching factor

observation points per public state $\leq (b+1)^k$

Number of last-infosets = "information asymmetry": measures *how much information is known to one player but not public to the team*

$k = 1$: team players have symmetric info

$\Rightarrow$ can be modeled as a perfect-recall player

tree size

CFR per-iteration runtime ∝ size of TB-DAG $\leq (b+1)^k \cdot \text{poly}(N)$

Polynomial time (and quite efficient in practice) when k is a constant!

Cannot be improved: $2^{o(k)} \cdot \text{poly}(N)$ would disprove ETH

How big is the TB-DAG?

CFR per-iteration runtime ∝ size of TB-DAG $\leq (b+1)^{k} \cdot \mathrm{poly}(N)$

What about a game like… contract bridge?

$k = 2 \cdot \binom{52}{13} < 2^{40}$, $b < 2^6$, $N \approx 2^{200}$ (very rough guess)

Naïve algorithm (write down the normal form and solve): time $\approx 2^{2^{200}}$

TB-DAG: $< 2^{6 \cdot 2^{40}} \cdot \mathrm{poly}(2^{200}) < 2^{2^{43}} \ll 2^{2^{200}}$

Bonus #1: Why timeability?

Partial answer:

Untimeable games of imperfect recall are weird…

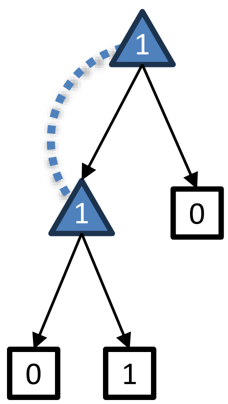

Every pure strategy scores 0
Randomizing uniformly and independently scores $1/4$

"Behavioral strategies $\not\subseteq$ mixed strategies?!"

Bonus #2: Hidden-role games

Luca Carminati*, Brian Hu Zhang*, Gabriele Farina, Nicola Gatti, Tuomas Sandholm
(EC 2024) “Hidden-Role Games: Equilibrium Concepts and Computation”

What Are Hidden-Role Games?

a.k.a. *social deduction* games

Adversarial team games in which *one team* ("Good") *does not know who its teammates are.*

Emphasis on *communication:* Players are free to talk to each other (e.g., to establish trust and coordinate actions), but Good players don't know whether they are talking to a Bad player!

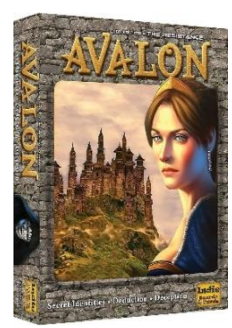

+ more applications: distributed systems, network security, federated learning, …

Highlights of Our Contributions

1. First **solution concepts** suitable for general hidden games

2. First **efficient algorithms** for solving in general hidden-role games

3. Application: **Exactly solve** 5- and 6-player versions of *Avalon*

Highlights of Our Contributions

1. First **solution concepts** suitable for general hidden games

2. First **efficient algorithms** for solving in general hidden-role games

3. Application: **Exactly solve** 5- and 6-player versions of *Avalon*

Solution Concepts for Hidden-Role Games

Can we just use…

- **Nash equilibrium?**
  - Most basic notion of equilibrium
  - Problem: Nash doesn't capture *team coordination.*

  (If there is only one team (Good or Bad), we want to capture the jointly-optimal strategy of the team, not simply a Nash equilibrium)

- **Team-correlated equilibrium?**
  - Most natural concept *when teams are common knowledge.*
  - Teammates can discuss strategy (incl. flipping random coins, i.e., correlating) *before the game begins*
  - Problem: in hidden-role games, before the game begins, you don't even know your teammates!

We need a new solution concept!

Solution Concepts for Hidden-Role Games

Solution: the *split-personality game*

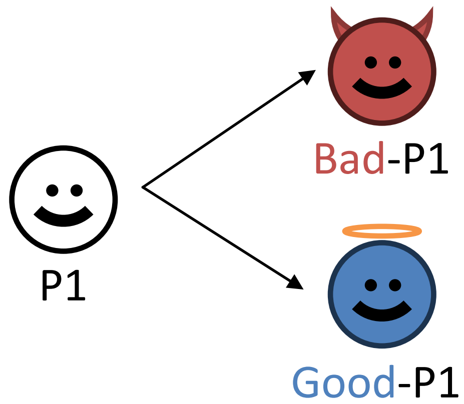

n-player **hidden-role game** → $2n$-player **adversarial team game**

Definition: Hidden-role equilibrium
= Team (uncorrelated) equilibrium of the split-personality game
= Optimal team strategy *for the good team* to commit to

Inherently asymmetric (no duality/minimax theorem)!

Asymmetry is *essential*: in any hidden-role game where the minimax theorem holds, Good can immediately cause teams to be publicly revealed by sending a long random string (→ no more hidden roles!)

Example

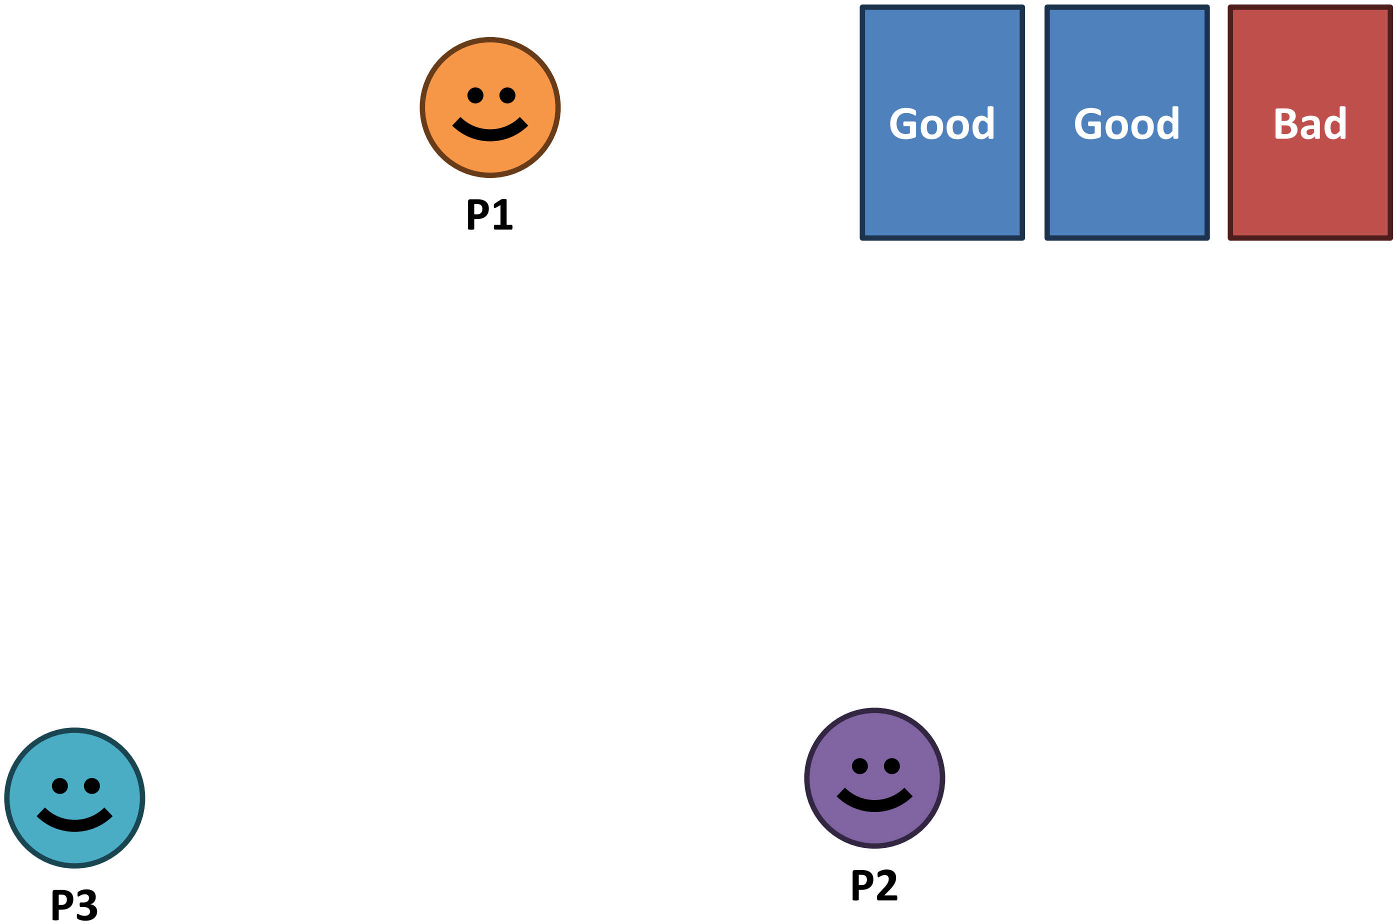

Example

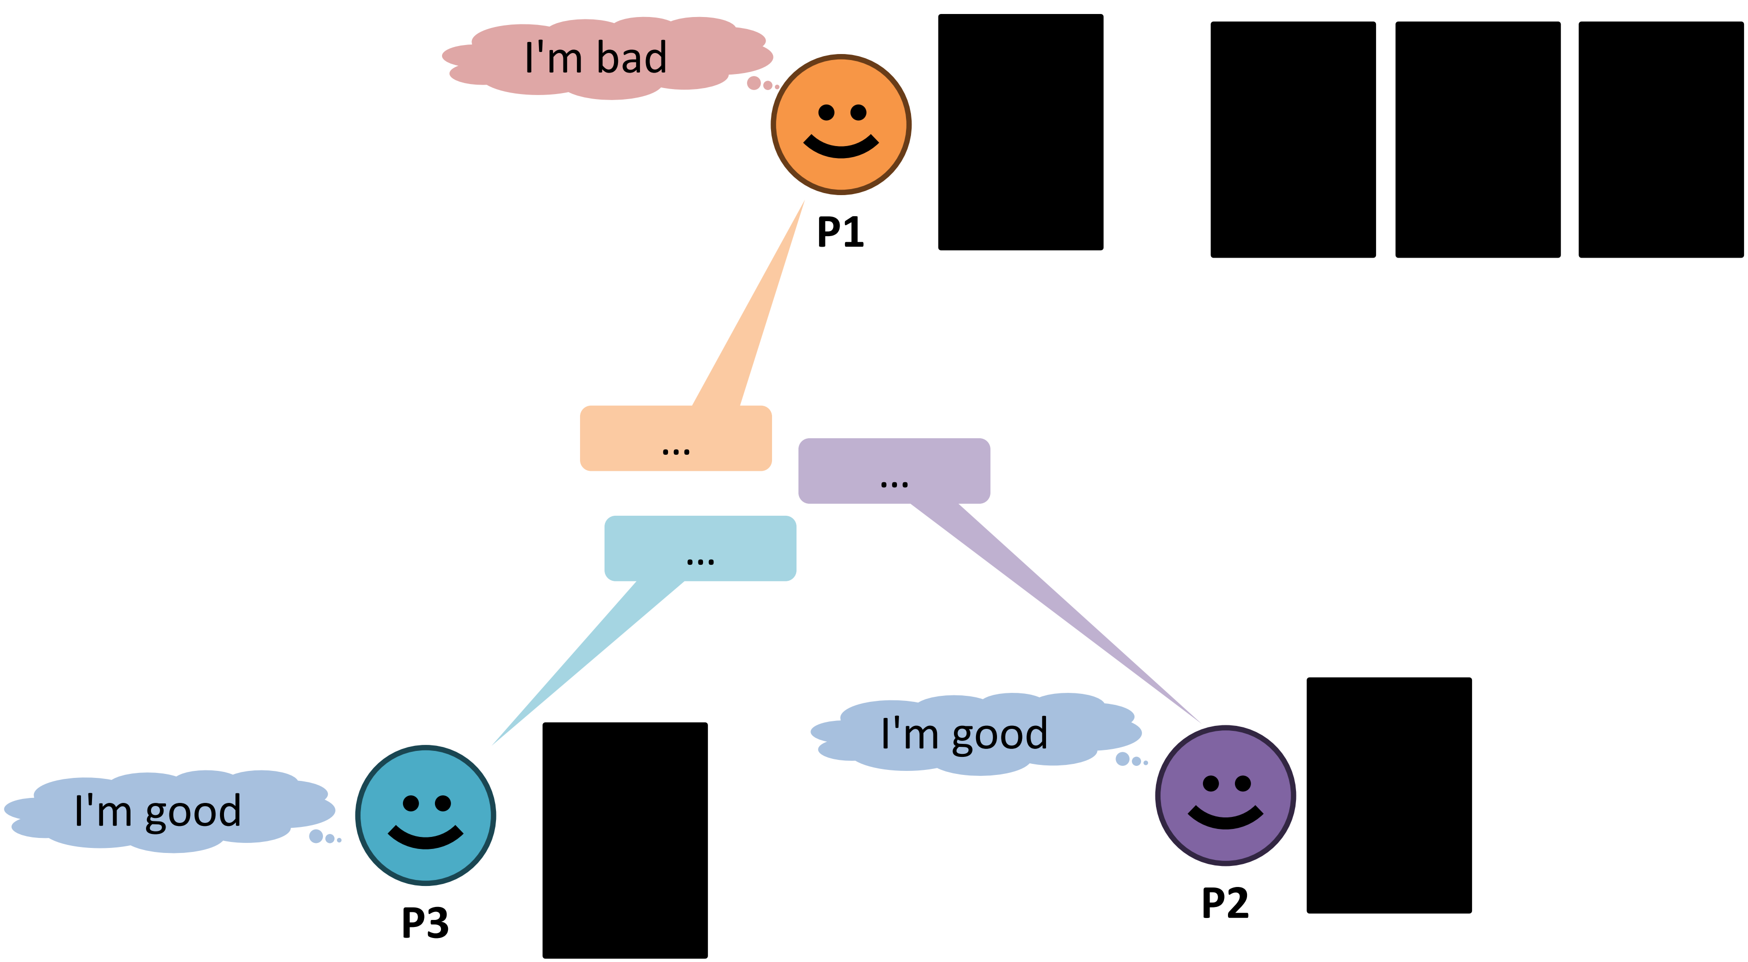

Example

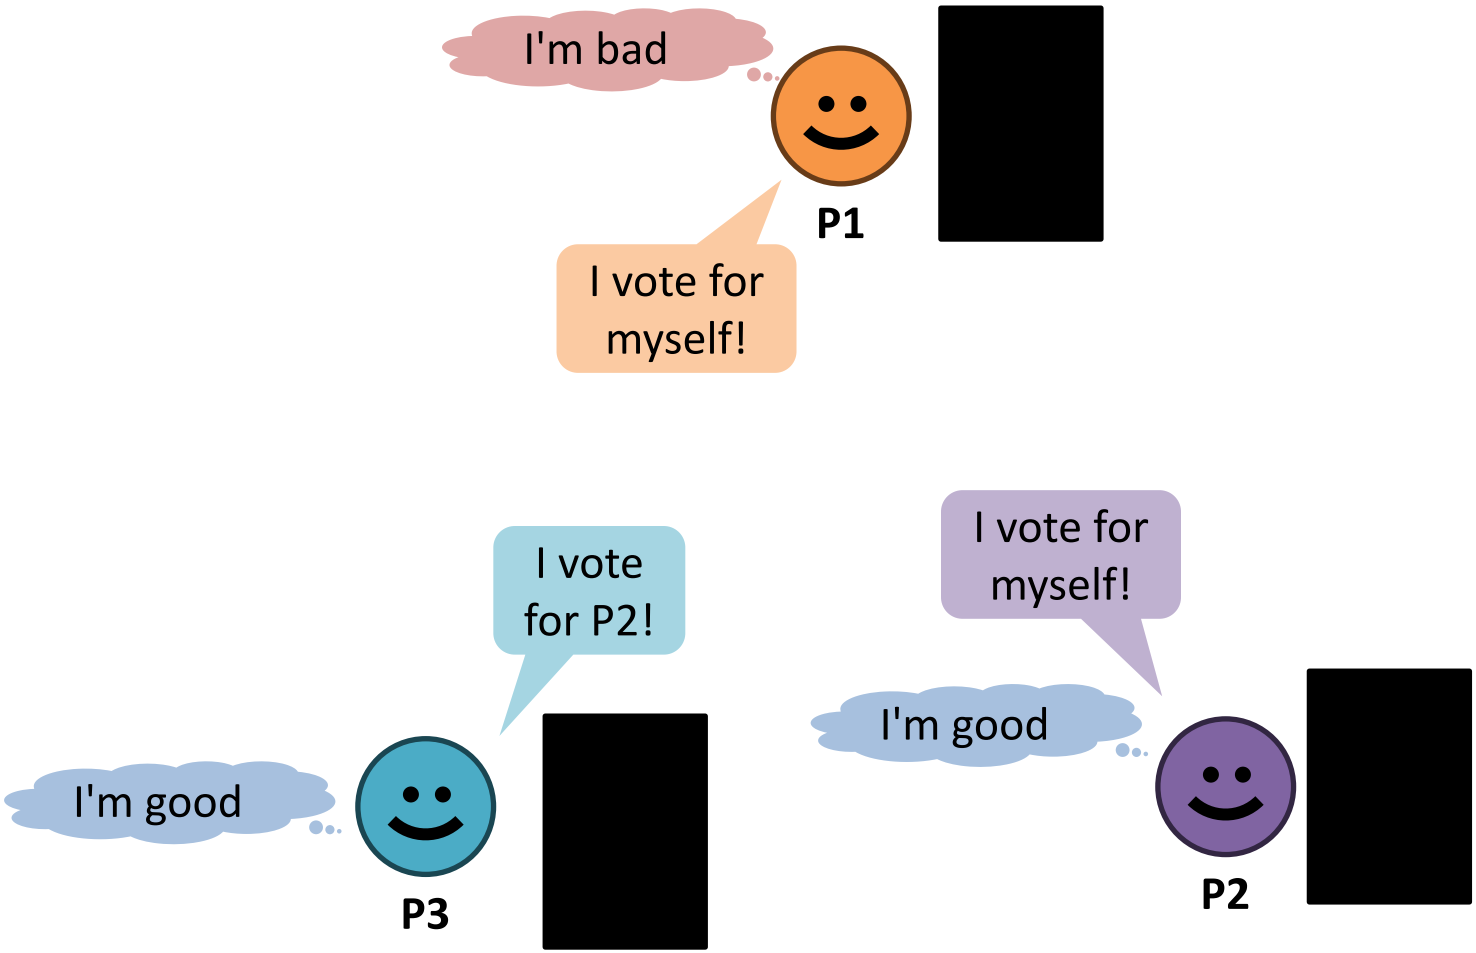
I'm bad
P1
I vote for myself!
I vote for P2!
I vote for myself!
I'm good
I'm good
P3
P2

Example

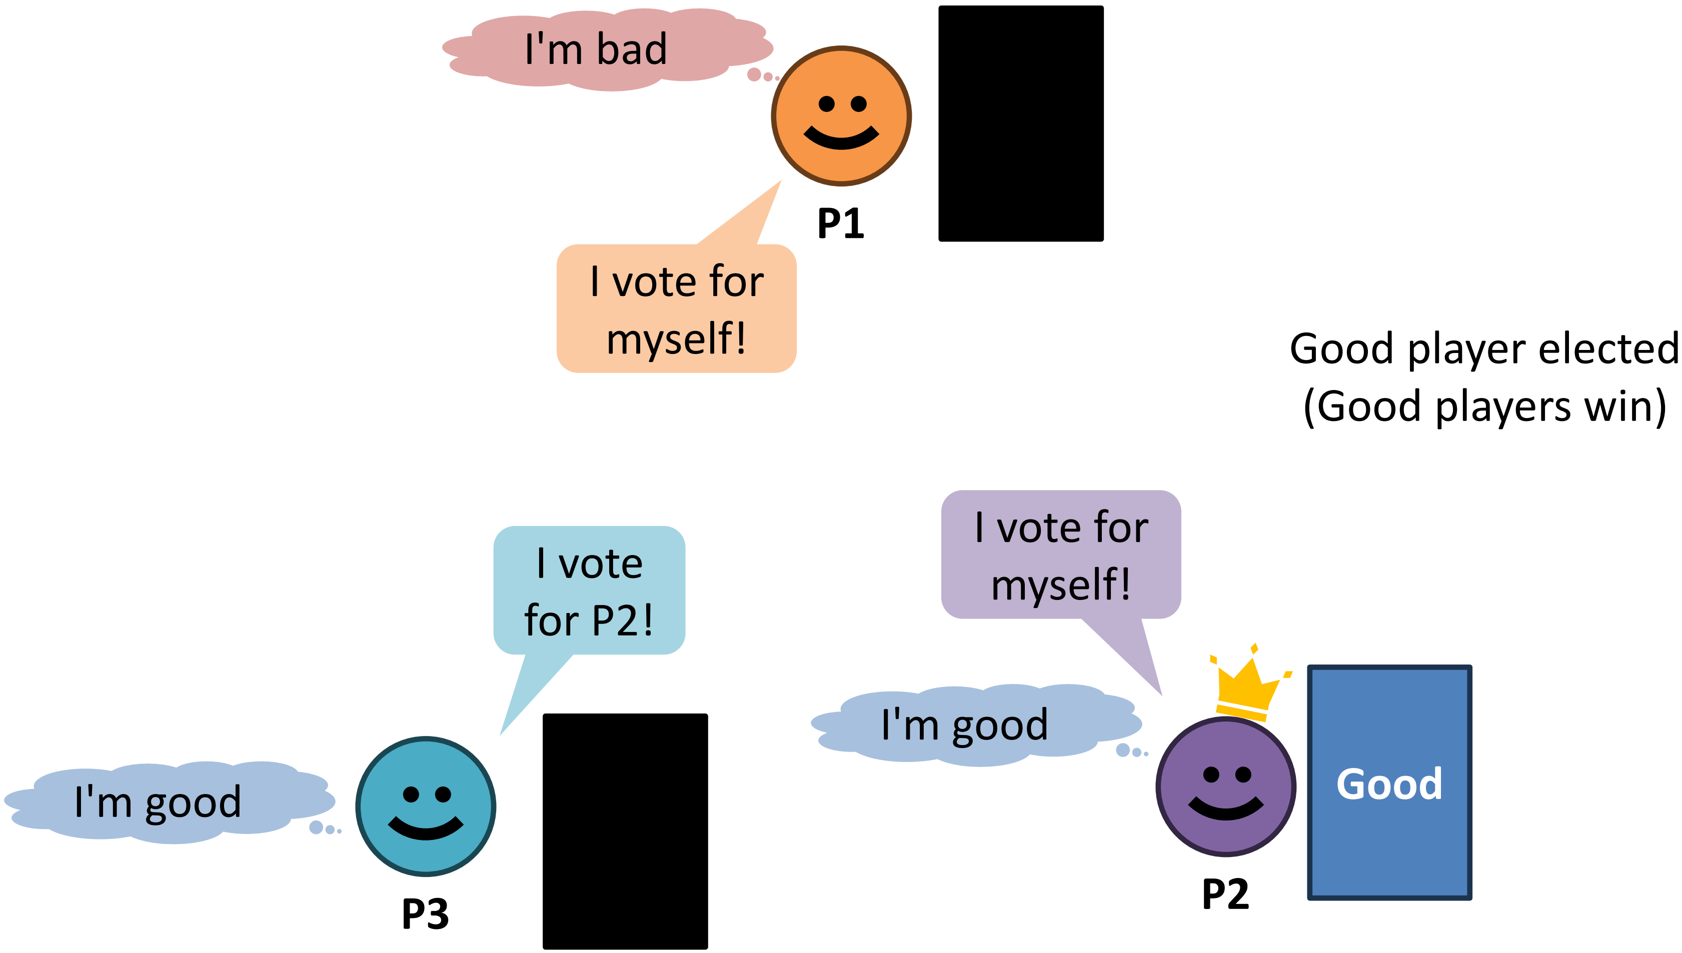
I'm bad
P1
I vote for myself!
Good player elected
(Good players win)
I vote for P2!
I'm good
P3
I vote for myself!
I'm good
Good
P2

Example

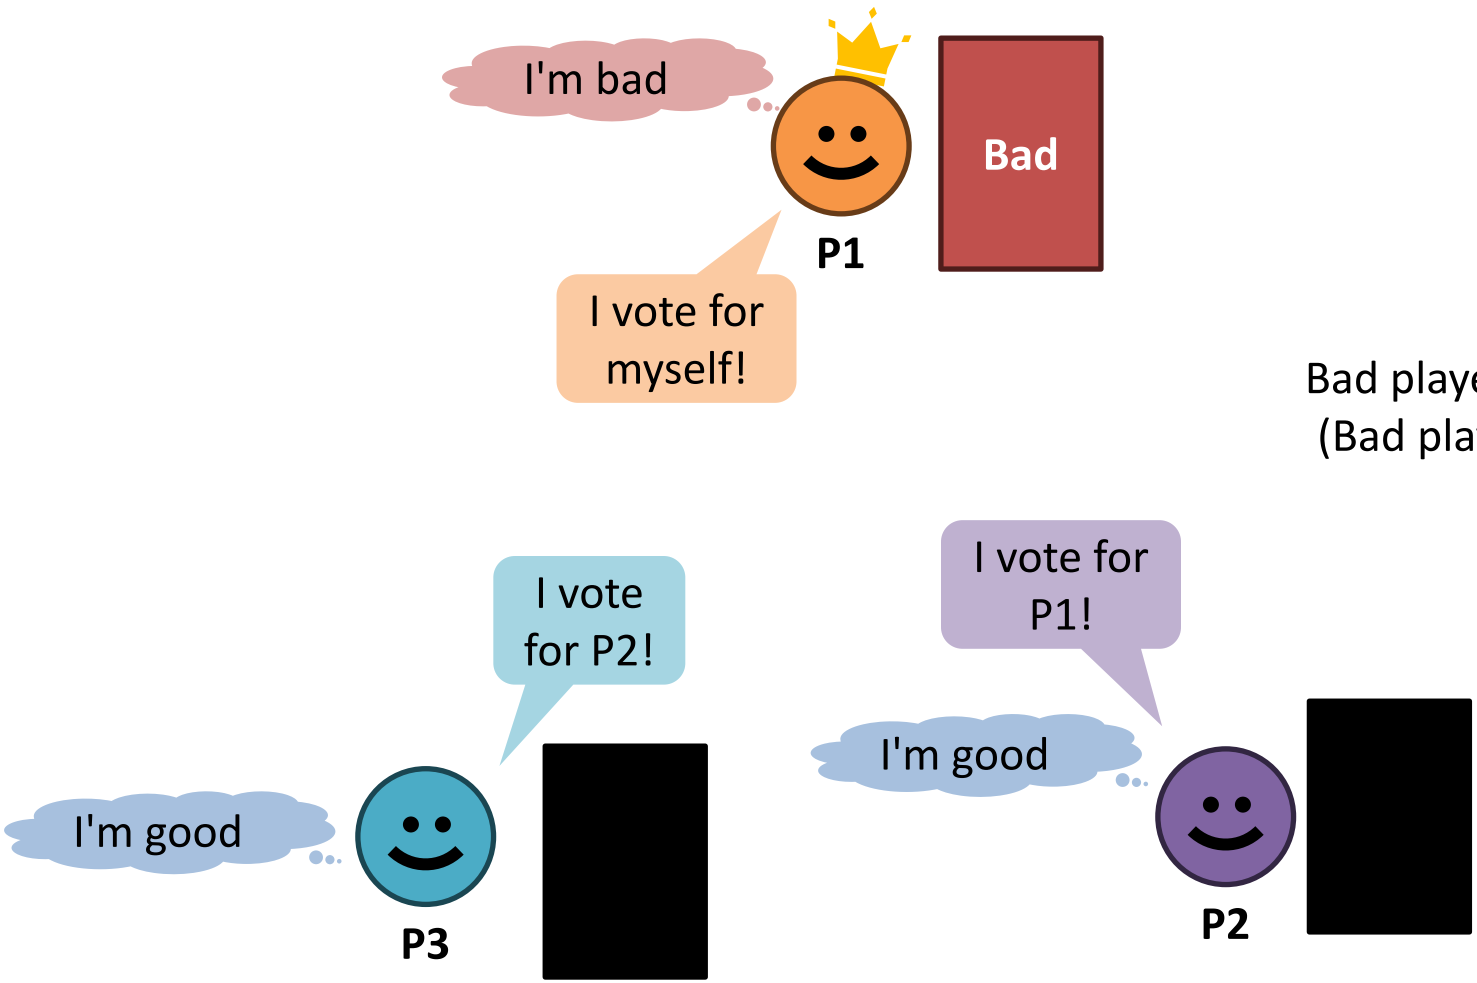

Bad player elected
(Bad player wins)

Example

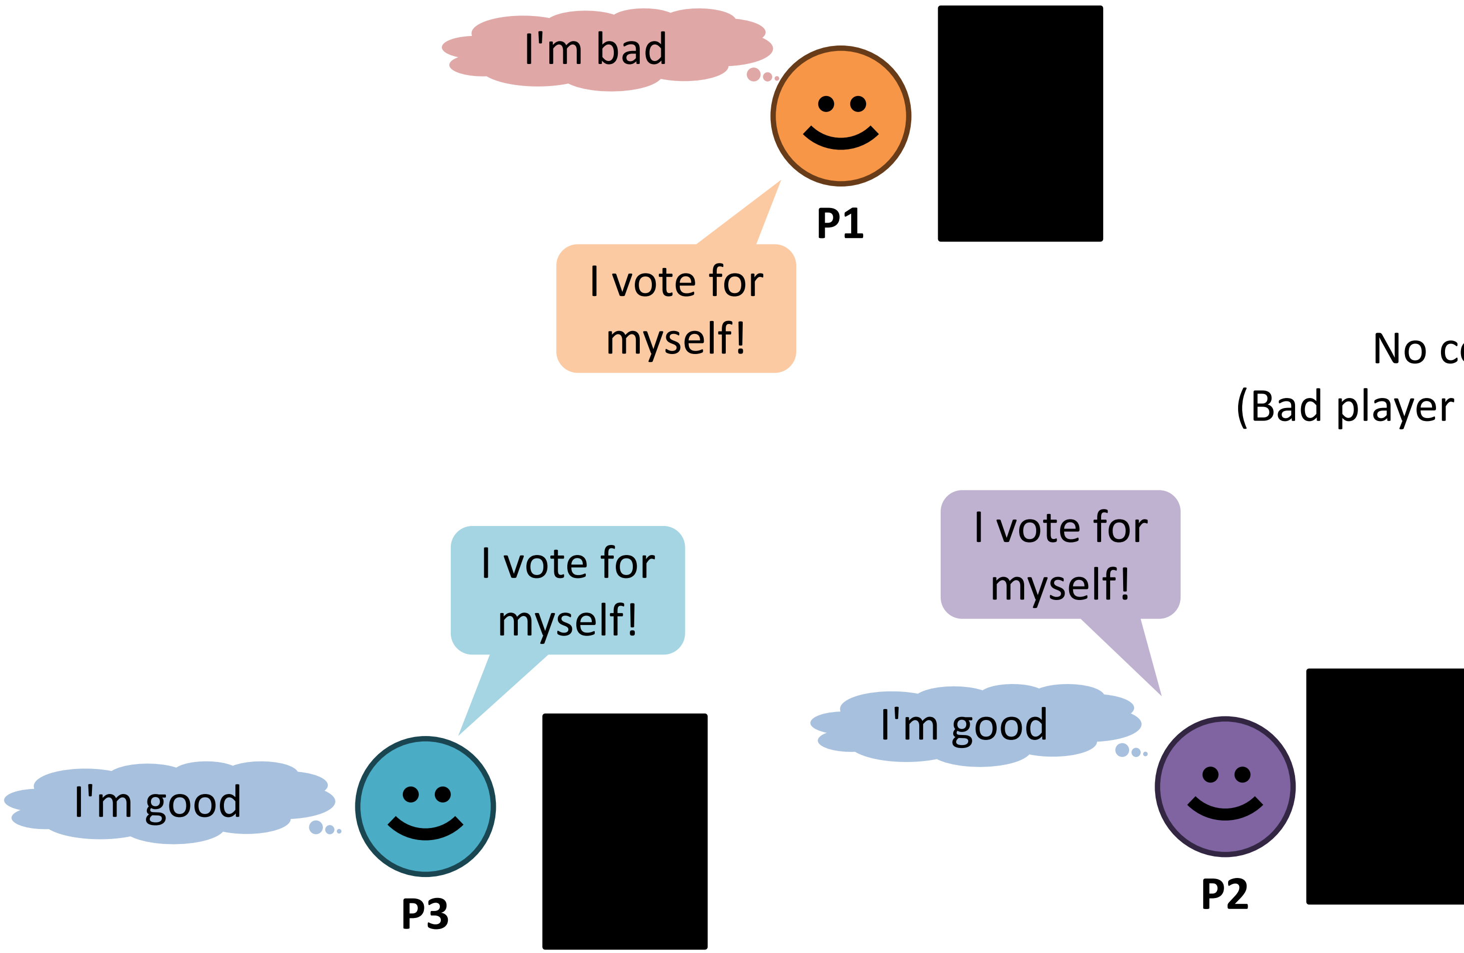

No consensus
(Bad player wins by default)

Example

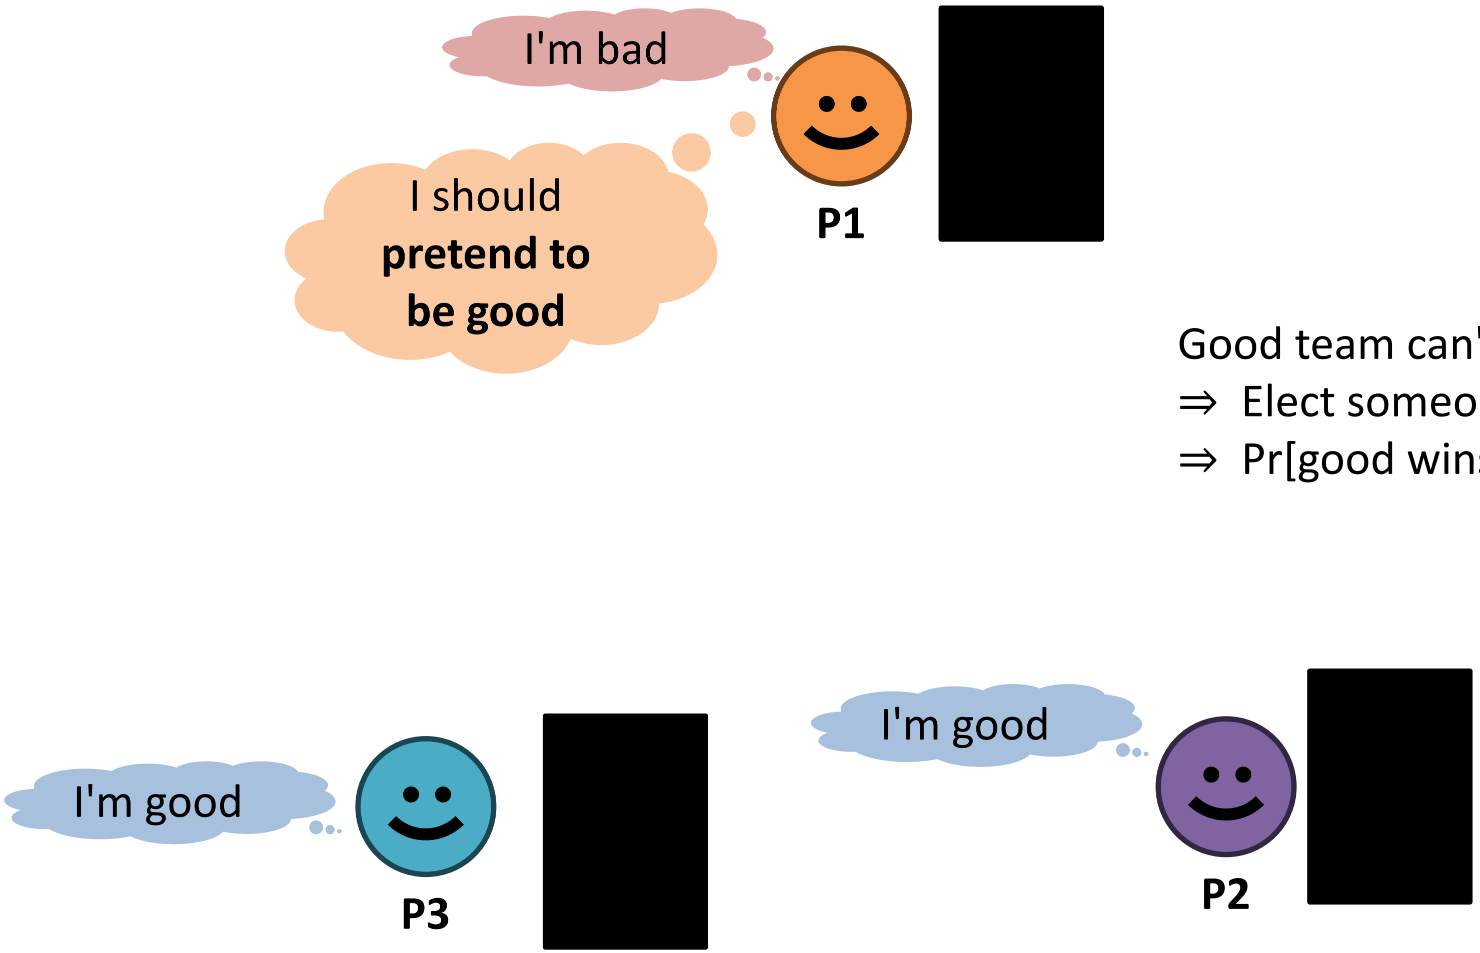

Good team can't learn anything
⇒ Elect someone arbitrary
⇒ Pr[good wins] = 2/3

Example

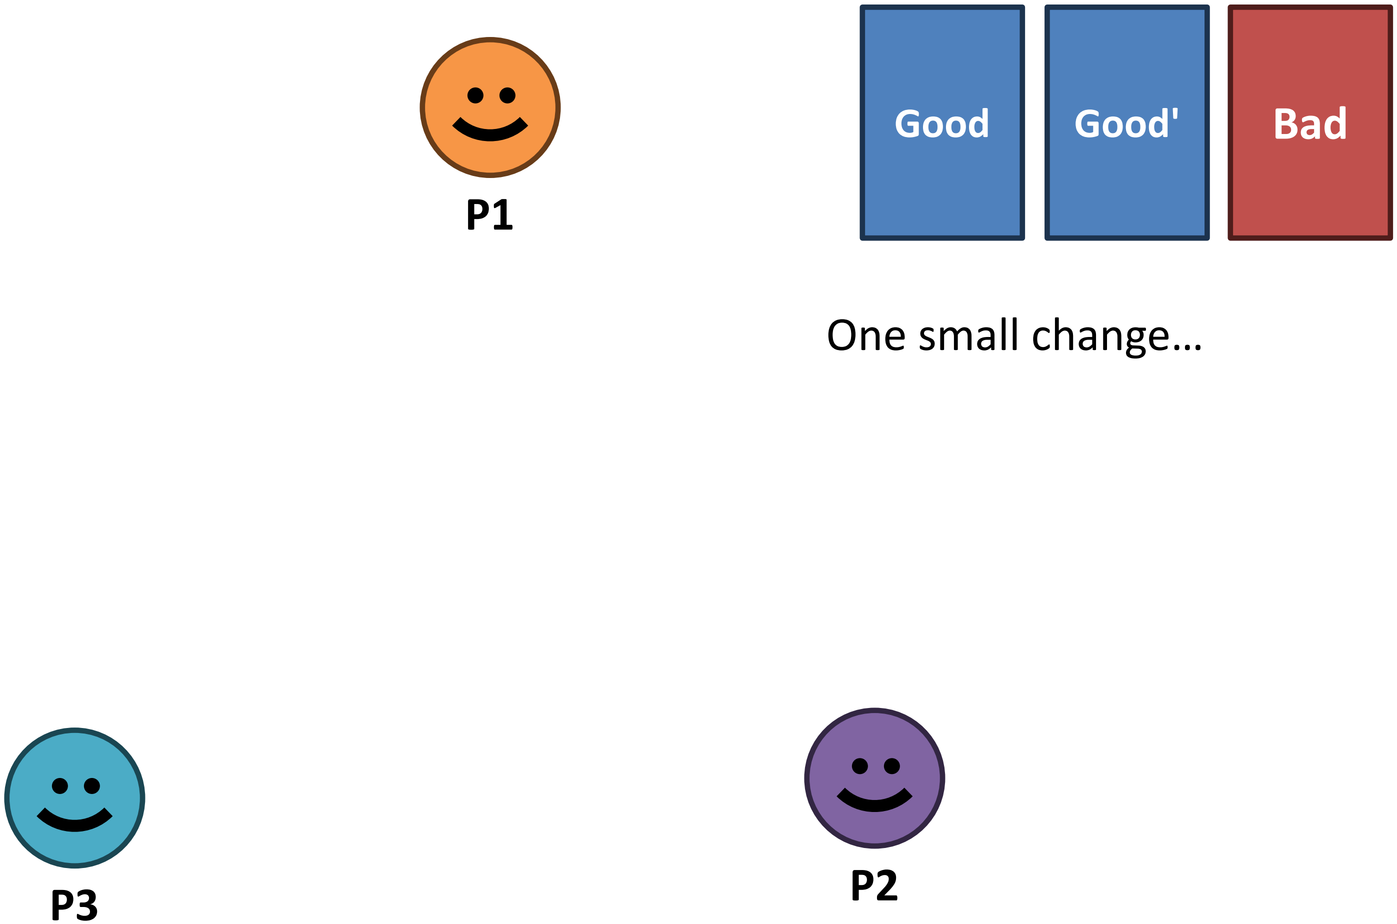

One small change...

Example

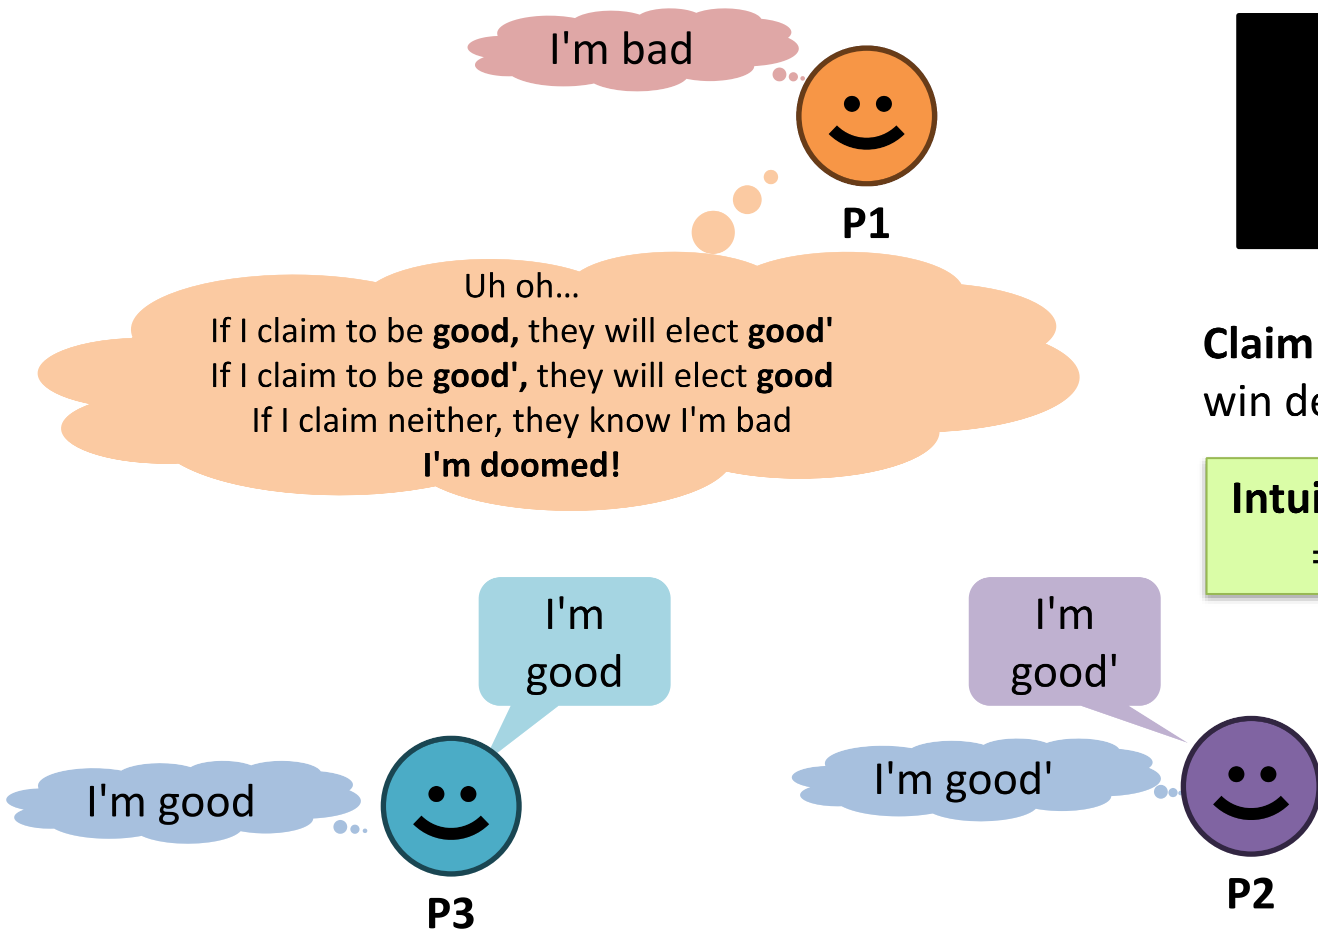

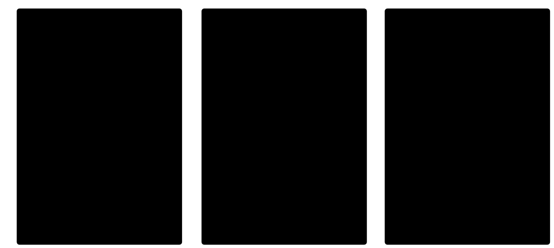

Claim: Good team can win deterministically

Intuition: Distinguishable cards ⇒ "implicit correlation"

Highlights of Our Contributions

1. First **solution concepts** suitable for general hidden games

2. First **efficient algorithms** for solving in general hidden-role games

3. Application: **Exactly solve** 5- and 6-player versions of *Avalon*

Computing Hidden-Role Equilibria

- Team Bad is a *minority* if it has fewer players than Good
- Team Bad is *coordinated* if (informally) the team "acts like it is being controlled by a single adversary controller"

Theorem [Main algorithmic result]:
There exists an **efficient algorithm** for computing hidden-role equilibria, assuming:
1. private communication is allowed,
2. Team Bad is a minority, and
3. Team Bad is coordinated,

Reasonable assumption: can be implemented with public communication + public-key crypto, assuming Bad is computationally bounded

Theorem [Lower bounds, informal]:
- Assuming (2) and (3) but not (1): NP-hard.
- Assuming (2) and (1) but not (3): coNP-hard.
- Assuming (2) only: Σ_2^P-hard.

Open problem: What about assuming (1) and (3)?

Proof Ideas for Main Theorem

Suppose that there is a player, who we call the *mediator*, who is *always* on the Good team.

Revelation principle:

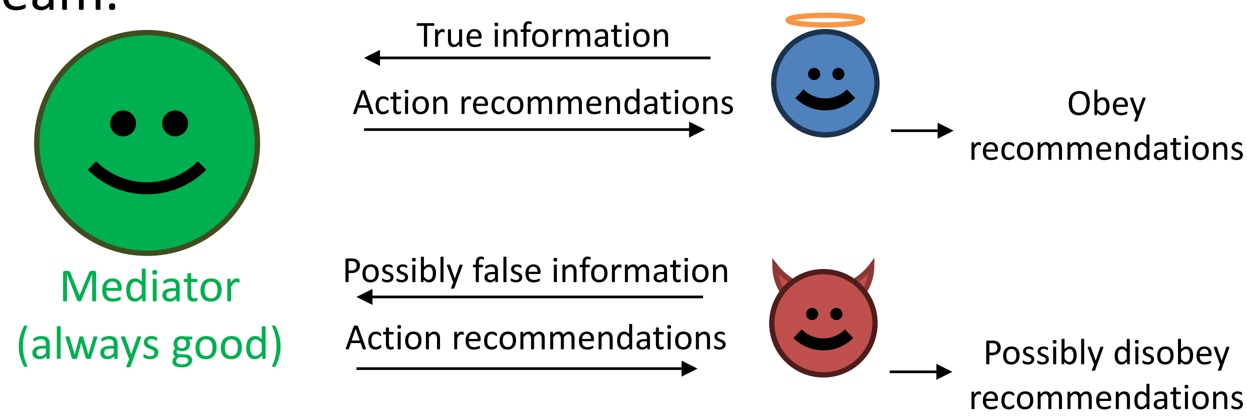

Main idea #1: With private communication, when Bad is *coordinated*, this is a **two-player zero-sum** game between the mediator and the Bad team

$\Rightarrow$ **can be solved efficiently!**

Implies that game values of hidden-role games are rational under these assumptions!

Main idea #2: If there is no mediator, **simulate one** using multi-party computation! (requires minority Bad team + private communication)

Highlights of Our Contributions

1. First **solution concepts** suitable for general hidden games

2. First **efficient algorithms** for solving in general hidden-role games

3. Application: **Exactly solve** 5- and 6-player versions of *Avalon*

Application: Exact solving of *Avalon*

- With < 7 players, the Bad team in *Avalon* is coordinated
  ⇒ main theorem applies!
- **Many** simplifications (e.g., removal of dominated strategies)
  ⇒ games small enough to solve exactly with LP

Variant	**5 Players**	**6 Players**
No special roles (*Resistance*)	$3 / 10 = 0.3000^*$	$1 / 3 \approx 0.3333^*$
Merlin	$2 / 3 \approx 0.6667^*$	$3 / 4 = 0.7500^*$
Merlin & Mordred	$731 / 1782 \approx 0.4102$	$6543 / 12464 \approx 0.5250$
Merlin & 2 Mordreds	$5 / 18 \approx 0.2778$	$99 / 340 \approx 0.2912$
Merlin, Mordred, Percival, Morgana	$67 / 120 \approx 0.5583$	—

*: Known to Christiano [2018]

Resistance $\cdot \frac{2}{3}$ = Merlin + 2 Mordreds?

No: distinguishable role cards for Good ⇒ implicit correlation!

Future Research

This is a **new class of games for which we can reasonably define what it means to "solve" a game!** Still many open questions & possible future directions:

- (*From earlier slide*) Efficient algorithm or hardness result for the case of coordinated, *non-minority* Bad team and private communication?
- Not even clear whether a *finite*-time alg exists if only public communication is allowed (how to bound the length of communication?)

  Especially important for practical interpretations, since humans playing these games often restrict ourselves to public communication + no crypto
- Other messaging structures? (e.g., *anonymous* messages)?
- What happens when *both* teams are hidden?

  Is there even a way to define hidden-role equilibria so that it does not depend on the seemingly-arbitrary choice of which team picks its strategy first?

Thank you!

References

Brian Hu Zhang, Gabriele Farina, Tuomas Sandholm (arXiv 2022; ICML 2023) "Team Belief DAG: Generalizing the Sequence Form to Team Games for Fast Computation of Correlated Team Max-Min Equilibria via Regret Minimization"

Stephen McAleer, Gabriele Farina, Gaoyue Zhou, Mingzhi Wang, Yaodong Yang, Tuomas Sandholm (NeurIPS 2023) "Team-PSRO for learning approximate TMECor in large team games via cooperative reinforcement learning"

Luca Carminati*, Brian Hu Zhang*, Gabriele Farina, Nicola Gatti, Tuomas Sandholm (EC 2024) "Hidden-Role Games: Equilibrium Concepts and Computation"