The Hardness of Signaling in Bayesian Games

Umang Bhaskar	**TIFR**
Yu Cheng	USC
Young Kun Ko	Princeton
Chaitanya Swamy	U Waterloo

Prisoners' Dilemma

	cooperate	defect
cooperate	-1 -1	0 -5
defect	-5 0	-4 -4

Prisoners' Dilemma

	cooperate	defect
cooperate	$-1+\theta$ / $-1+\theta$	0 / $-5+\theta$
defect	$-5+\theta$ / 0	-4 / -4

$$\theta \sim U\{-2,-1,0,1,2\}$$

- no information: (D,D) is NE
- reveals θ:

  (C,C) is NE if $\theta \geq 1$ (w.p. 2/5)

  (D,D) is NE o.w. (w.p. 3/5)

[example modified from Dughmi '14]

Prisoners' Dilemma

	cooperate	defect
cooperate	$-1+\theta$ (blue), $-1+\theta$ (red)	0 (blue), $-5+\theta$ (red)
defect	$-5+\theta$ (blue), 0 (red)	-4 (blue), -4 (red)

$$\theta \sim U\{-2,-1,0,1,2\}$$

- **H** if $\theta \geq 0$, **L** otherwise

  H: (C,C) is NE (w.p. 3/5)

  L: (D,D) is NE (w.p. 2/5)

(C,C) is NE w.p. 3/5!

[example modified from Dughmi '14]

Bayesian Game

- payoffs are uncertain
- depend on state of nature θ

θ_1	…	θ_M

w.p. λ_1 … λ_M

(λ prior, known)

e.g., $\theta \sim U\{-2,-1,0,1,2\}$

Principal knows θ,

chooses Signaling Scheme:

States $\rightarrow$ Signals

$\Theta \rightarrow \Sigma$ (possibly randomized)

$-2,\quad -1,\quad 0,\quad 1,\quad 2$

$-2, -1 \rightarrow$ L

$0, 1, 2 \rightarrow$ H

Bayesian Game

- payoffs are uncertain
- depend on state of nature θ

θ_1	…	θ_M

w.p. λ_1 … λ_M

(λ prior, known)

e.g., $\theta \sim U\{-2,-1,0,1,2\}$

- scheme 1: same signal for all θ
  - reveals nothing beyond prior
  - called **no revelation**
- scheme 2: diff signal for each θ
  - reveals θ
  - called **full revelation**

running time poly(# states, size of game)

Problem: What is computational complexity of optimal signaling scheme?

(e.g., maximize social welfare, player 1’s payoff, etc.)

Problem: What is computational complexity of *optimal* signaling scheme?

for 2-person zero-sum games

(goal: maximize row-player's payoff)

Problem: What is computational complexity of *optimal* signaling scheme?

for 2-person zero-sum games

(goal: maximize row-player's payoff)

since

- NE (essentially) unique,
- well-understood,
- poly-time computable

Approximation

A **signaling scheme** for instance I is ϵ-approximate
if R's payoff is $\geq OPT(I) - \epsilon$

Algorithm A is ϵ-approximate if on any instance I,
it computes an ϵ-approximate signaling scheme.

Algorithm A is an **FPTAS** if, given $\epsilon > 0$,
computes ϵ-approximate signalling scheme
in time poly$\left(\frac{1}{\epsilon}\right)$.

Previous Work for 0-Sum Games

- Design problem first studied by Dughmi
- Obtaining an FPTAS is as hard as
  recovering a planted clique in a random graph [Dughmi '14]

Previous Work for 0-Sum Games

- Design problem first studied by Dughmi
- Obtaining an FPTAS is as hard as recovering a planted clique in a random graph [Dughmi '14]
- Given $\epsilon > 0$, ϵ-approximate signalling scheme can be computed in time poly$(n^{\log n/\epsilon^2})$ [CCDEHT '14]

Previous Work for 0-Sum Games

Independently,

- Obtaining an FPTAS is NP-hard
- For a constant $\epsilon > 0$,
  computing poly-time ϵ-approximate signalling scheme
  is as hard as constructing sub-exponential time algo for SAT

[Rubinstein '15]

Results for 2-player 0-sum games

Result I: NP-hard to obtain an FPTAS

$$\left(\begin{array}{l}\text{algorithm that given } \epsilon > 0, \\ \text{computes } \epsilon\text{-approximate signalling scheme} \\ \text{in time poly}\left(\frac{1}{\epsilon}\right).\end{array}\right)$$

Results for 2-player 0-sum games

Result II: For a constant $\epsilon > 0$,
computing ϵ-approximate signalling scheme
is as hard as recovering a planted clique in a random graph.

- why not NP-hard?

$\exists$ quasi-polynomial time algorithm for this [CCDEHT '15]

so NP-hardness would give a QPT algo for an NP-hard problem

$n^{O(\log n)}$

Results for 2-player 0-sum games

Result II: For a constant $\epsilon > 0$,
computing ϵ-approximate signalling scheme
is as hard as recovering a planted clique in a random graph.

- why not NP-hard?
- planted-clique hardness:
  - n vertices, each edge exists w.p. ½

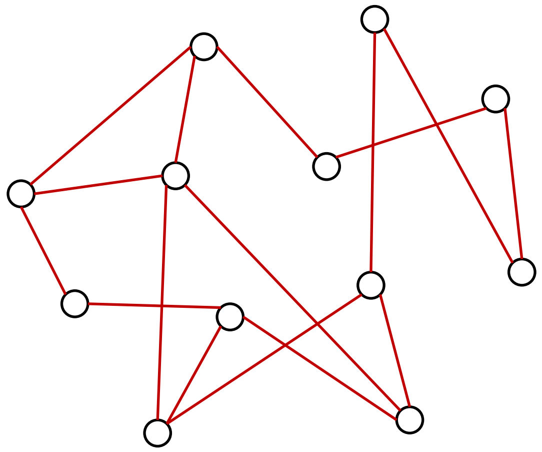

Results for 2-player 0-sum games

Result II: For a constant $\epsilon > 0$,
computing ϵ-approximate signalling scheme
is as hard as recovering a planted clique in a random graph.

- why not NP-hard?
- planted-clique hardness:
  - n vertices, each edge exists w.p. ½
  - select k random vertices, create k-clique

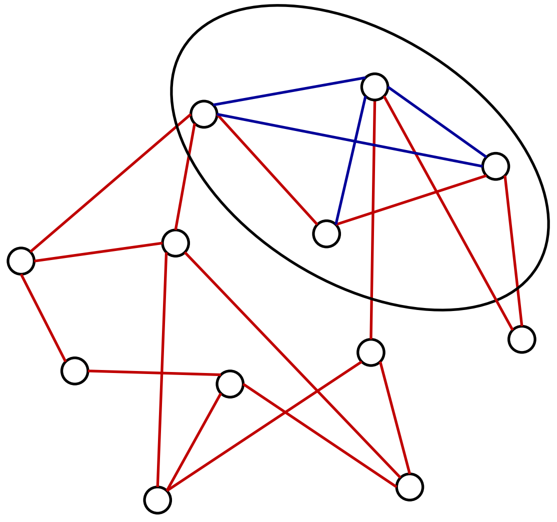

Problem: find planted clique

This Talk

Theorem: NP-hard to compute optimal signalling scheme

- reduction from Balanced Complete Bipartite Subgraph (BCBS) problem…
- but to the dual separation problem

This Talk

Theorem: NP-hard to compute optimal signalling scheme

Step 1: Signaling is at least as hard as threshold signaling, the separation problem for the dual

Step 2: Threshold signaling problem is NP-hard via reduction from BCBS

Posteriors

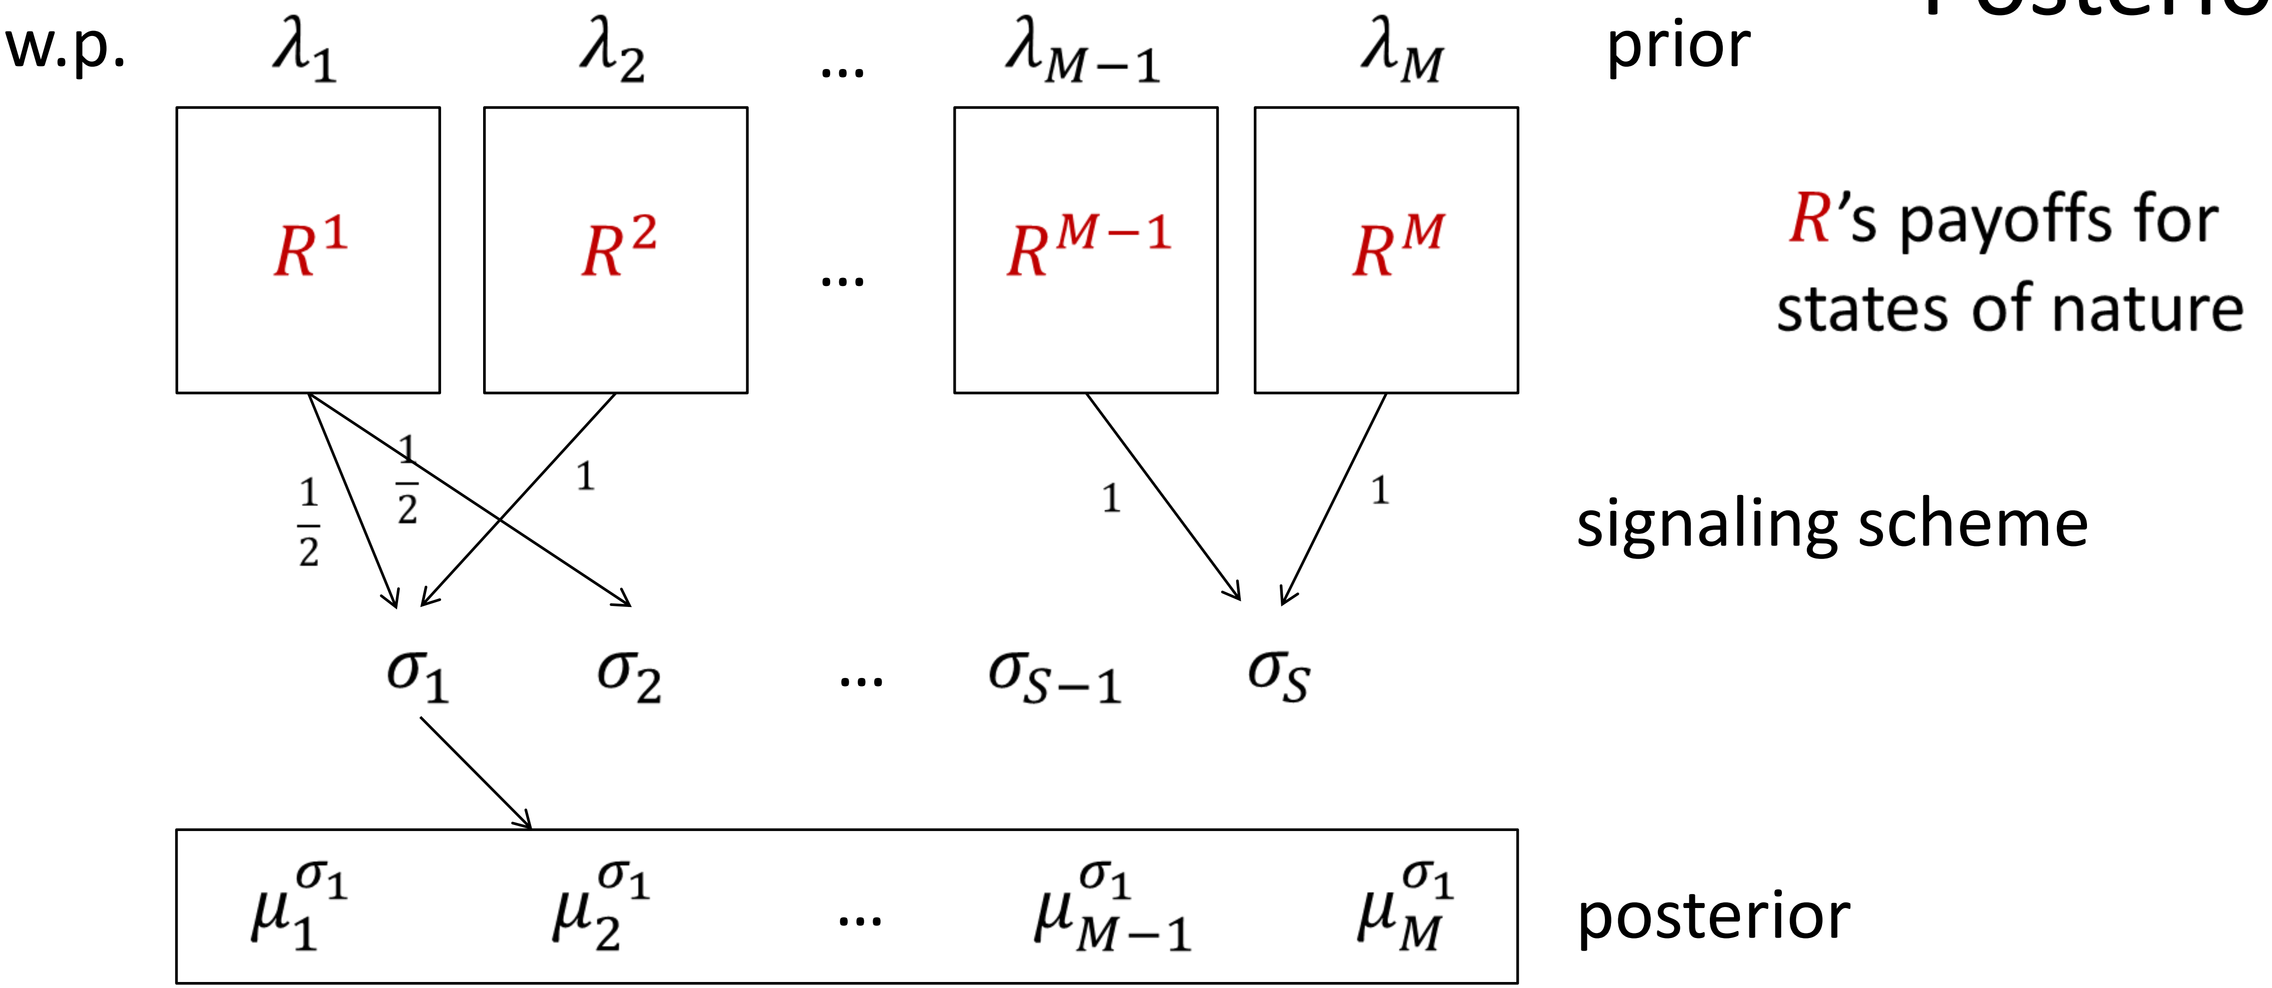

$$\mathrm{E}[R|\sigma_1] \quad = \quad \mu_1^{\sigma_1} R^1 + \mu_2^{\sigma_1} R^2 + \ \ldots + \ \mu_M^{\sigma_1} R^M$$

Posteriors

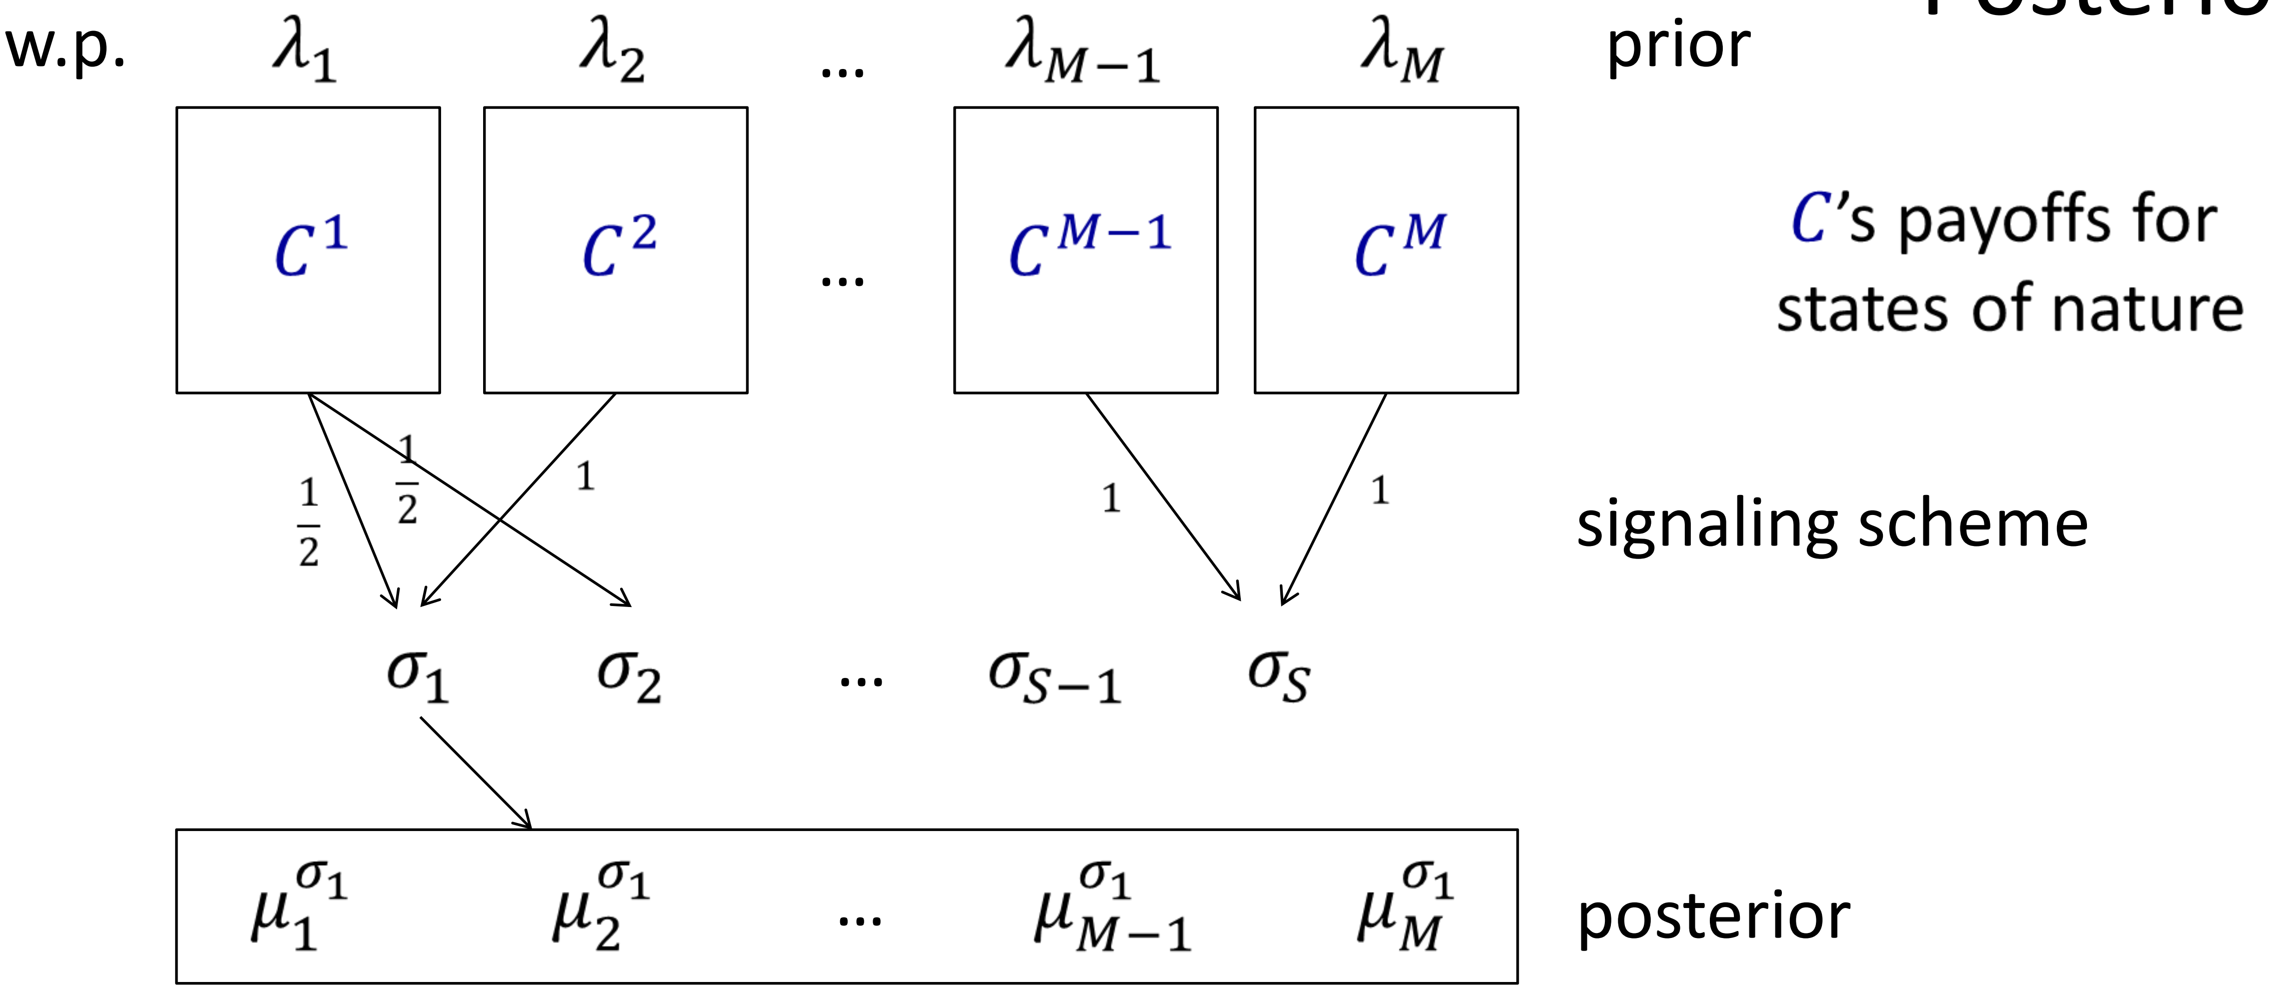

$$\mathrm{E}[C|\sigma_1] = \mu_1^{\sigma_1} C^1 + \mu_2^{\sigma_1} C^2 + \ldots + \mu_M^{\sigma_1} C^M$$

An Example

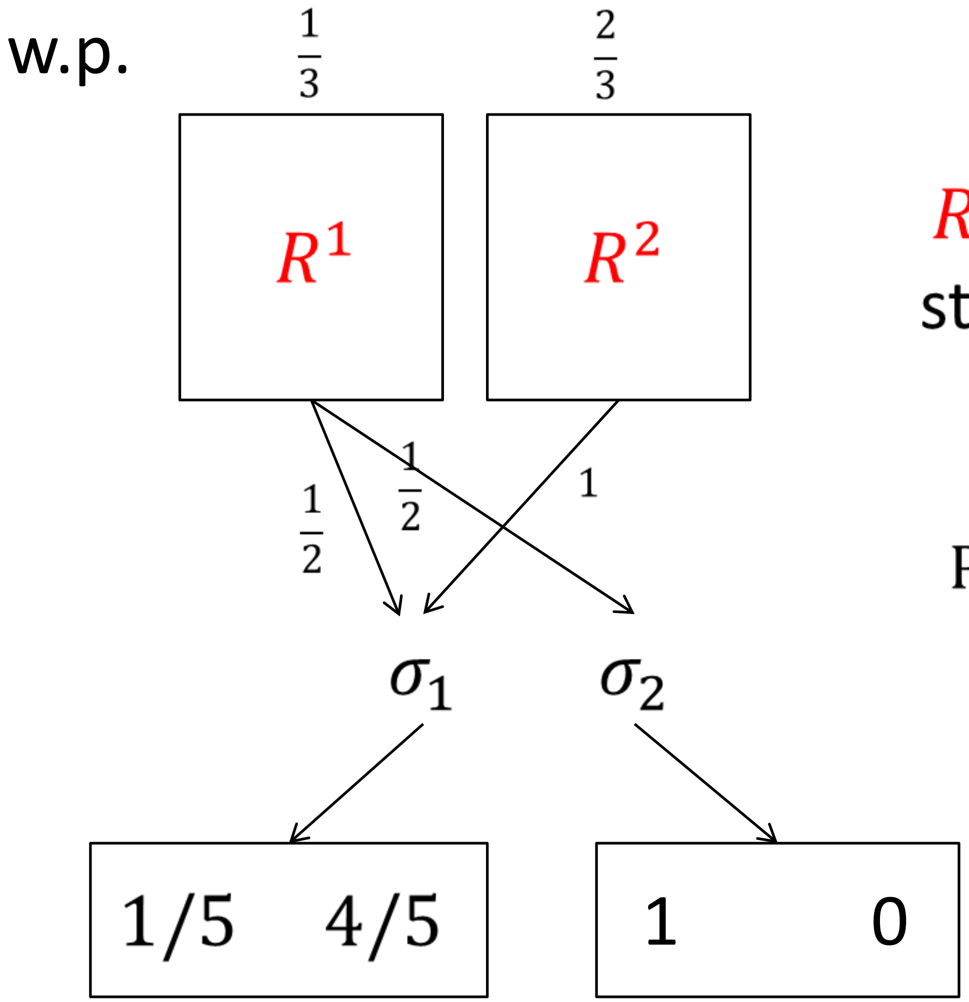

prior

R's payoffs for states of nature

$$\Pr[\sigma_1] = 1/6 + 2/3 = 5/6 \qquad \Pr[\sigma_2] = 1/6 + 0 = 1/6$$

posterior

$$\Pr[\sigma_1] \times \boxed{\Pr[\theta_1|\sigma_1]} + \Pr[\sigma_2] \times \boxed{\Pr[\theta_1|\sigma_2]} = \Pr[\theta_1]$$

Posterior Distributions

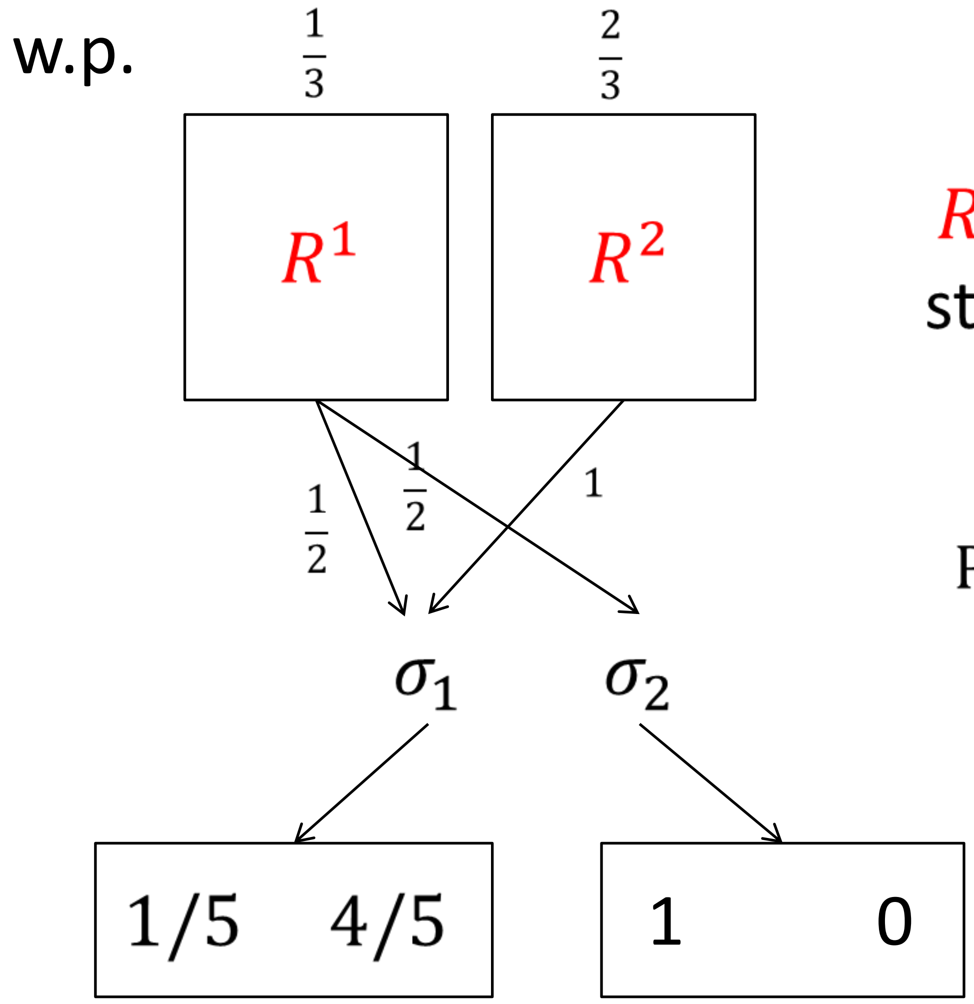

prior

R's payoffs for states of nature

$$\Pr[\sigma_1] = 1/6 + 2/3 = 5/6$$

$$\Pr[\sigma_2] = 1/6 + 0 = 1/6$$

- each signal σ gives a posterior distr over states of nature
- posteriors form convex decomposition of the prior

Posterior Distributions

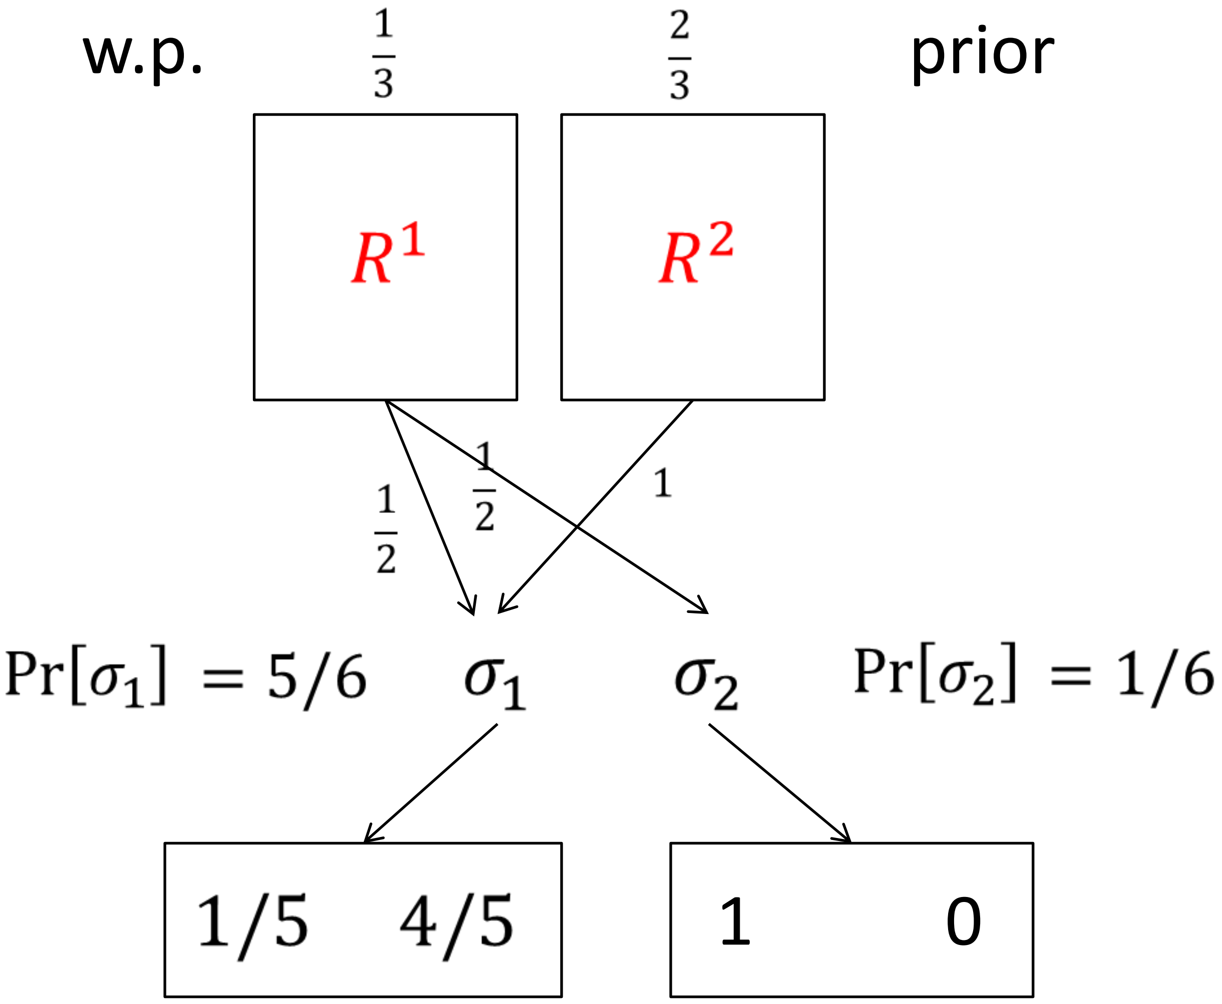

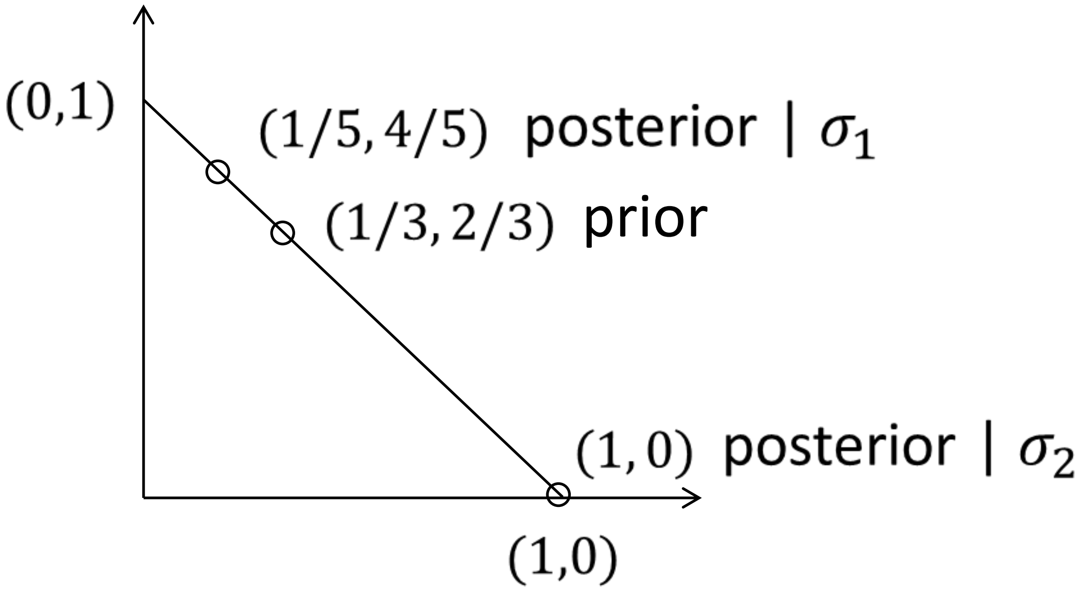

prior = $\Pr[\sigma_1] \times$ posterior | σ_1

$+ \Pr[\sigma_2] \times$ posterior | σ_2

Posterior Distributions

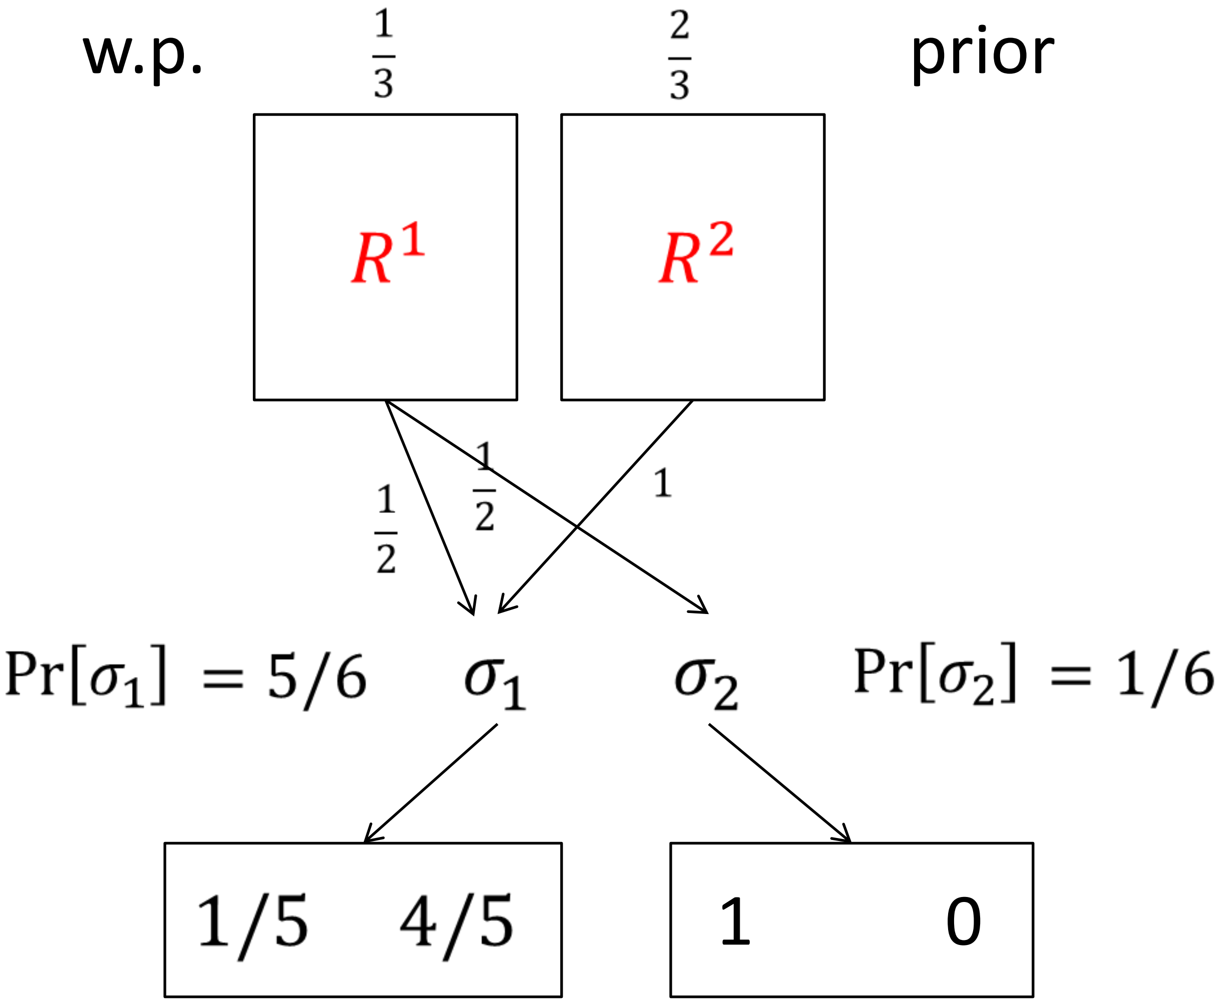

We shift focus and try to find **posteriors** rather than signals.

Conditions:

- each posterior is distribution over states of nature
- posteriors form convex decomposition of prior

Computing Optimal Posteriors

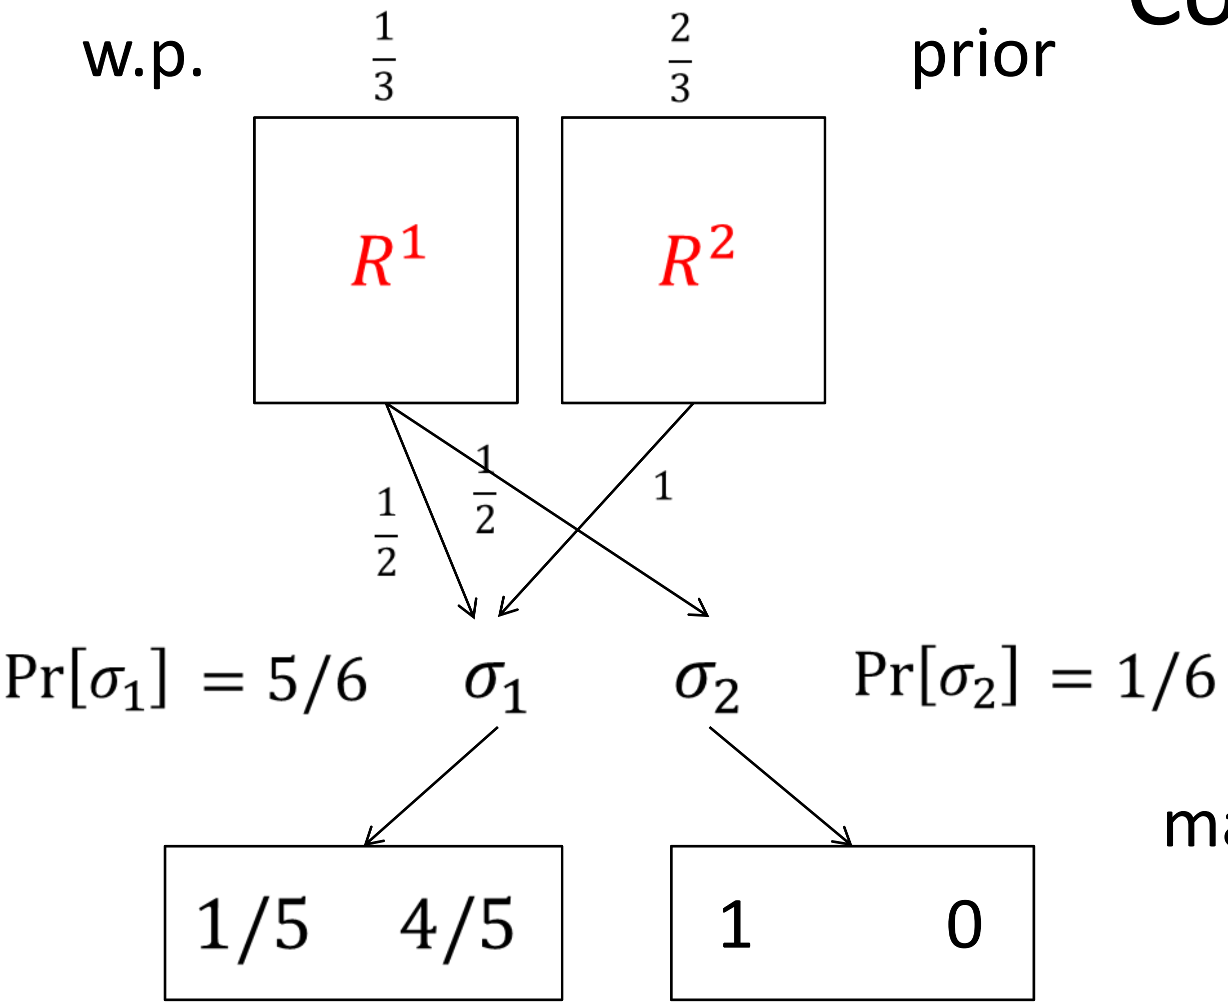

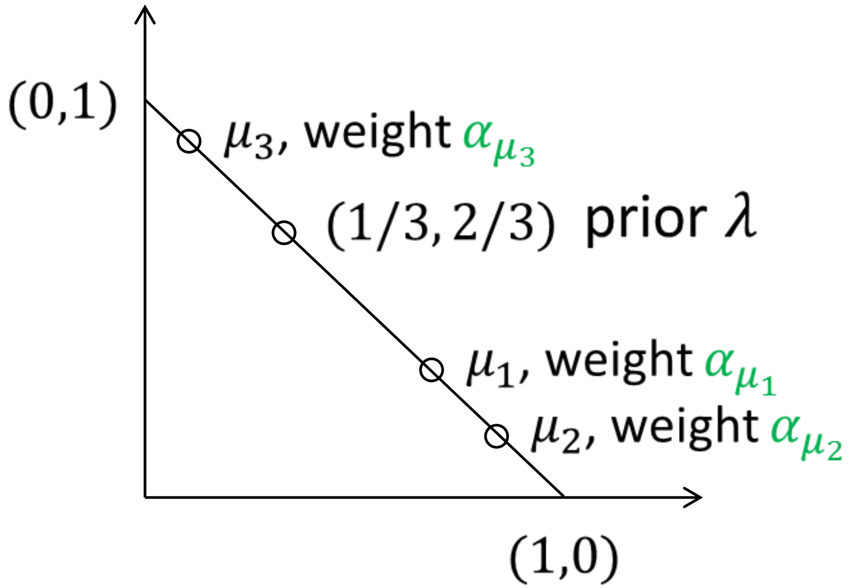

max $\sum_\mu \alpha_\mu \times R$'s payoff with posterior μ

$$\alpha_{\mu_1}\mu_1 + \alpha_{\mu_2}\mu_2 + \alpha_{\mu_3}\mu_3 + \cdots = \lambda$$

$$\alpha_{\mu_1} + \alpha_{\mu_2} + \alpha_{\mu_3} + \cdots = 1, \quad \alpha_\mu \geq 0$$

Computing Optimal Posteriors

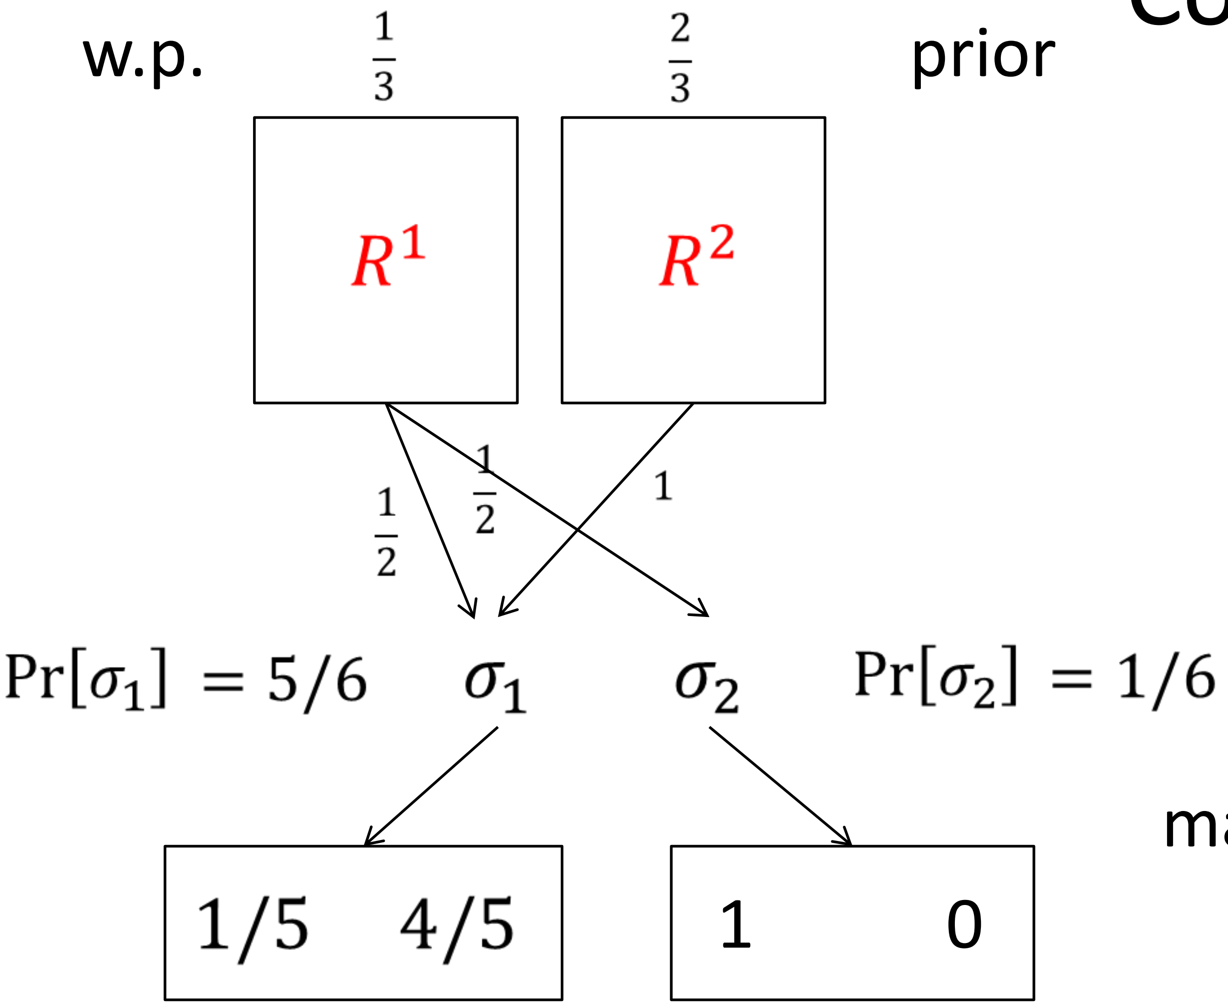

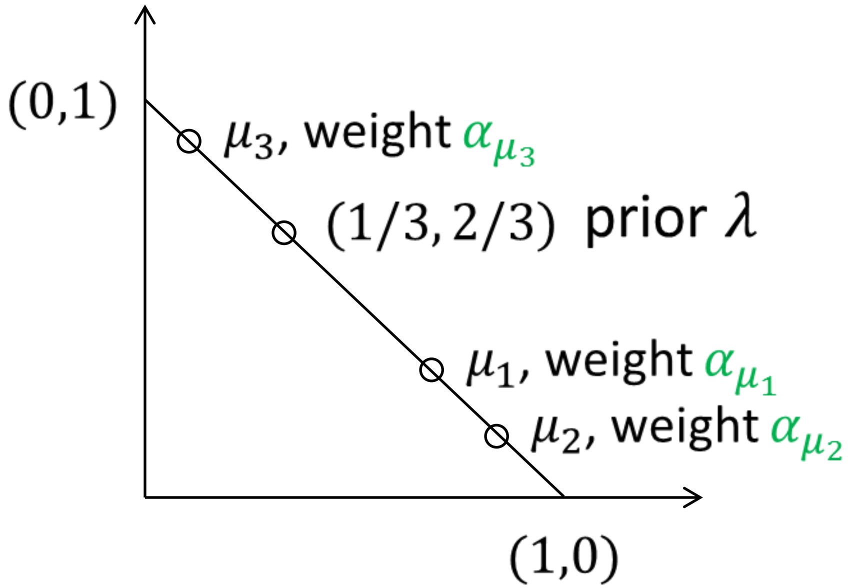

$$\max \ \sum_\mu \alpha_\mu \times \mathrm{val}(\mu)$$

$$\alpha_{\mu_1}\mu_1 + \alpha_{\mu_2}\mu_2 + \alpha_{\mu_3}\mu_3 + \cdots = \lambda$$

$$\alpha_{\mu_1} + \alpha_{\mu_2} + \alpha_{\mu_3} + \cdots = 1, \quad \alpha_\mu \geq 0$$

A Linear Program for Signaling

- this is a linear program

$$\max \sum_{\mu \in \Delta_M} \alpha_\mu \; \mathrm{val}(\mu)$$

$$\text{s.t.} \sum_{\mu \in \Delta_M} \alpha_\mu \; \mu \; = \; \lambda$$

$$\sum_{\mu \in \Delta_M} \alpha_\mu = 1, \;\; \alpha_\mu \geq 0$$

(implied by previous constraint)

Δ_M: set of **all** distributions over states

(note: infinite set, hence infinite variables)

Dual Linear Program for Signaling

$$\max \sum_{\mu \in \Delta_M} \alpha_\mu \ \mathrm{val}(\mu)$$

$$\text{s.t.} \sum_{\mu \in \Delta_M} \alpha_\mu \ \mu = \lambda$$

$$\alpha_\mu \geq 0$$

Δ_M: set of **all** distributions over states

- this is a linear program
- consider the dual linear program:

$$\min \ w^T \lambda$$

$$w^T \mu \geq \mathrm{val}(\mu) \quad \text{for all } \mu \in \Delta_M$$

The Separation Problem

$$\min \quad w^T \lambda$$

$$w^T \mu \geq \text{val}(\mu) \quad \text{for all } \mu \in \Delta_M$$

λ: prior distribution over states

Δ_M: set of **all** distributions over states

$\text{val}(\mu)$: R's payoff with posterior μ

Theorem: Solving an LP is as hard as finding a violated constraint, given variable values

(optimization $\equiv$ separation)

Thus, solving dual LP

$\equiv$

given w, $\exists ? \ \mu : w^T \mu < \text{val}(\mu)$

Threshold Signaling Problem

$$\min \quad w^T \lambda$$

$$w^T \mu \geq \text{val}(\mu) \quad \text{for all } \mu \in \Delta_M$$

λ: prior distribution over states

Δ_M: set of **all** distributions over states

$\text{val}(\mu)$: R's payoff with posterior μ

Thus, solving dual LP

$$\equiv$$

given w, $\exists? \ \mu : w^T\mu < \text{val}(\mu)$

A simpler problem:

given $c \in \mathrm{R}$, $\exists? \ \mu : c < \text{val}(\mu)$

This is the Threshold Signaling Problem

Threshold Signaling Problem

$$\min \quad w^T \lambda$$

$$w^T \mu \geq \text{val}(\mu) \quad \text{for all } \mu \in \Delta_M$$

λ: prior distribution over states

Δ_M: set of **all** distributions over states

$\text{val}(\mu)$: R's payoff with posterior μ

Thus, solving dual LP

$$\equiv$$

given w, $\exists?\ \mu : w^T\mu < \text{val}(\mu)$

A simpler problem:

given $c \in \mathrm{R}$, $\exists?\ \mu : c < \text{val}(\mu)$

Theorem: **Signaling** is at least as hard as given $c \in \mathrm{R}$, $\exists?\ \mu : c < \text{val}(\mu)$

Threshold Signaling Problem

Theorem: **Signaling** is at least as hard as given $c \in \mathrm{R}$, $\exists?\ \mu : c < \mathrm{val}(\mu)$

μ: distribution over states, $\mathrm{val}(\mu)$: R's payoff with posterior μ

Proof: (1) strong LP duality,

(2) optimization $\equiv$ separation

but need to consider infinite-dimensionality, etc.

(details skipped)

Threshold Signaling Problem

Theorem: **Signaling** is at least as hard as given $c \in \mathbb{R}$, $\exists ?\ \mu : c < \mathrm{val}(\mu)$

μ: distribution over states, $\mathrm{val}(\mu)$: R's payoff with posterior μ

Step 1: Signaling is at least as hard as threshold signaling

Step 2: Threshold signaling problem is NP-hard

reduce [Balanced Complete Bipartite Subgraph (BCBS)] to threshold signaling in [Network Security Games]

Balanced Complete Bipartite Subgraph Problem

Balanced Complete Bipartite Subgraph Problem

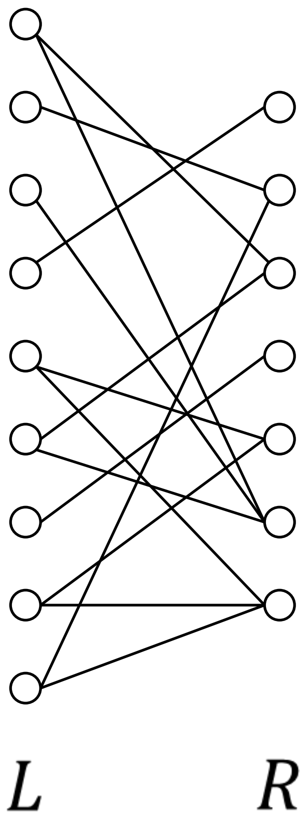

Given bipartite graph $G = (L \cup R, E)$, integer r,

does it contain $K_{r,r}$, complete bipartite graph of size r?

Balanced Complete Bipartite Subgraph Problem

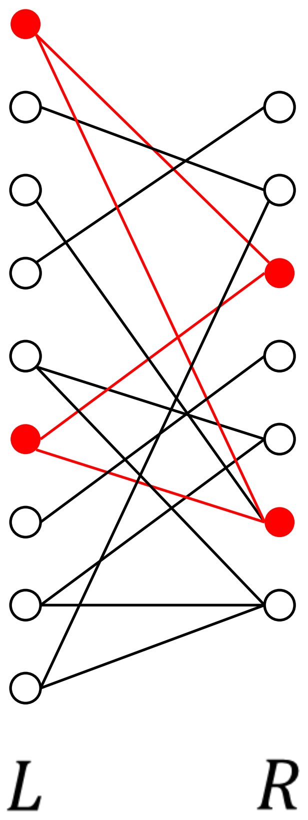

$K_{2,2}$, no $K_{3,3}$

Given bipartite graph $G = (L \cup R, E)$, integer r,
does it contain $K_{r,r}$, complete bipartite graph of size r?

Theorem: BCBS is NP-complete [GJ '79]

Network Security Games

Network Security Games

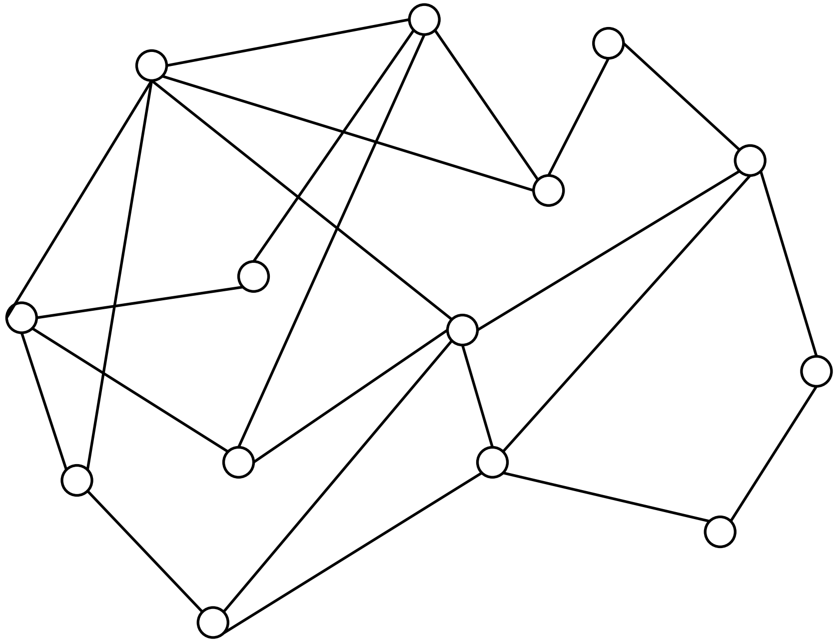

Given $G = (V, E)$,

V = states of nature

= strategies of R

= strategies of C

R's payoff: +1 if R adjacent to θ

- 1 if $C = R$ or $C = \theta$

(R must defend θ from C)

Network Security Games

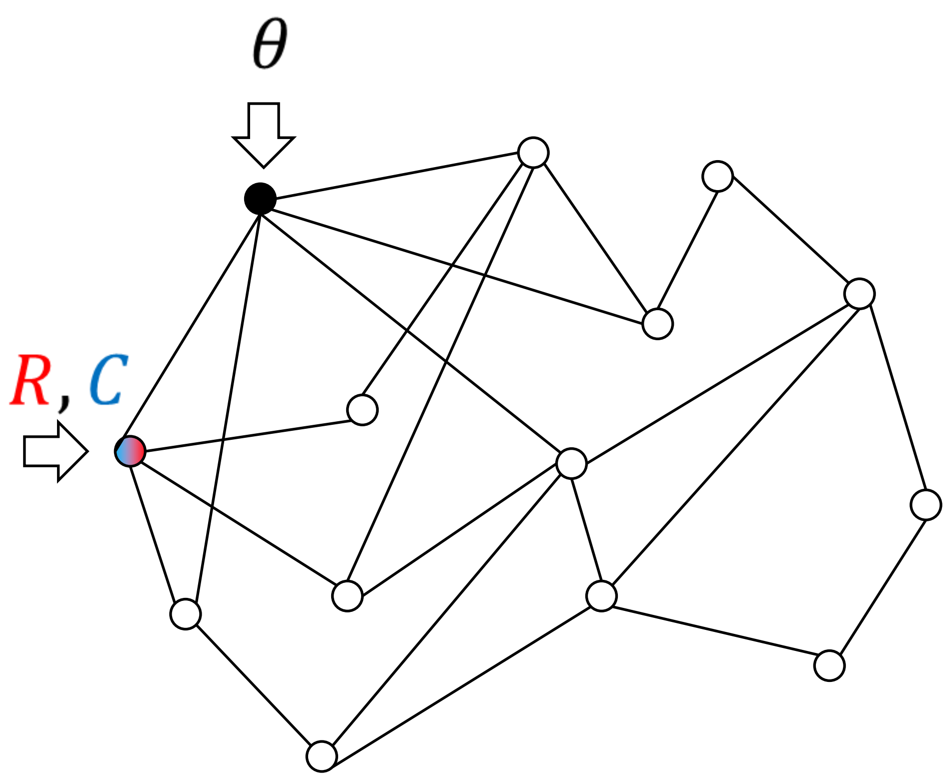

R's payoff = 1 – 1 = 0

Given $G = (V, E)$,

V = states of nature

= strategies of R

= strategies of C

R's payoff: +1 if R adjacent to θ

- 1 if $C = R$ or $C = \theta$

(R must defend θ from C)

Network Security Games

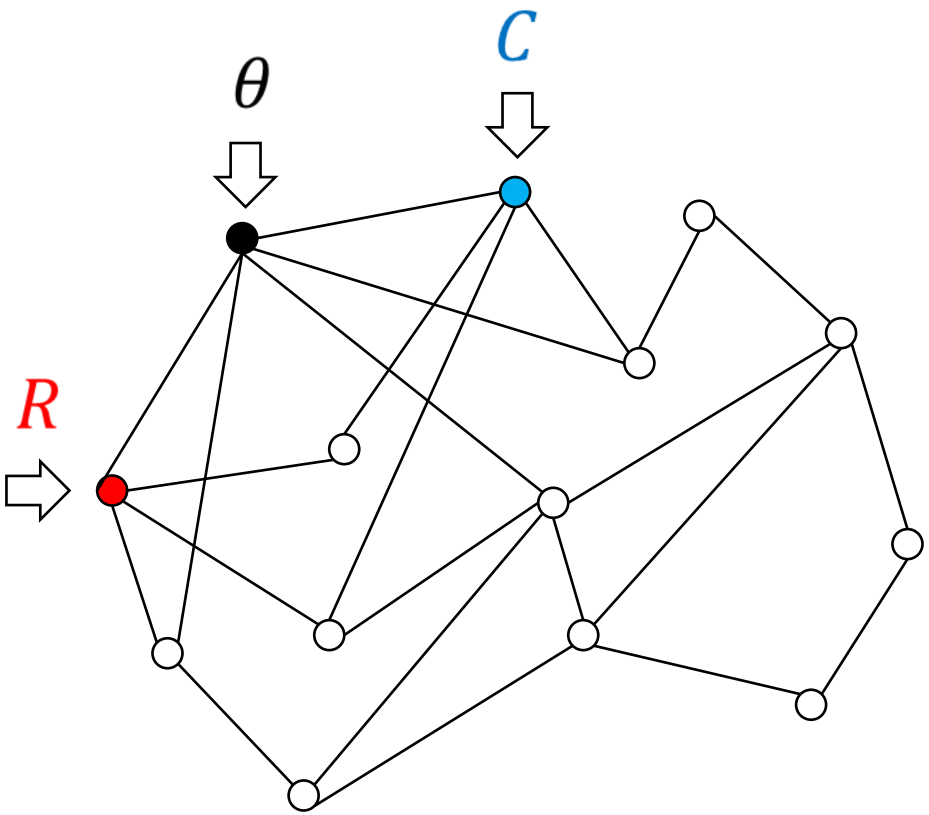

R's payoff = 1

Given $G = (V, E)$,

V = states of nature
= strategies of R
= strategies of C

R's payoff: +1 if R adjacent to θ
- 1 if $C = R$ or $C = \theta$

Network Security Games

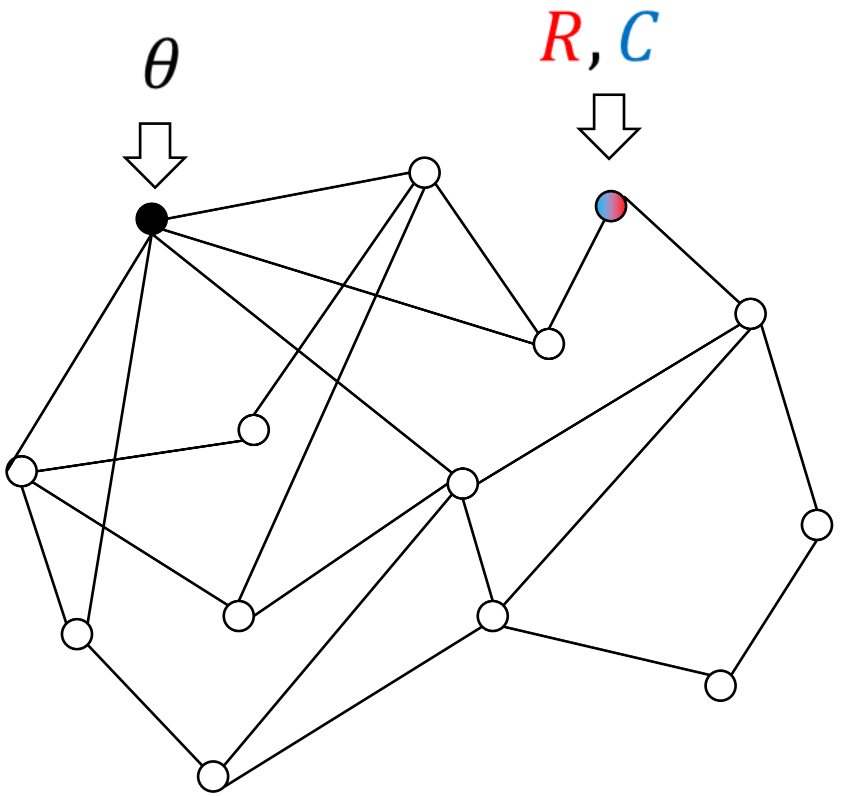

Given $G = (V, E)$,

V = states of nature

= strategies of R

= strategies of C

R's payoff: +1 if R adjacent to θ

- 1 if $C = R$ or $C = \theta$

R's payoff = -1

Network Security Games

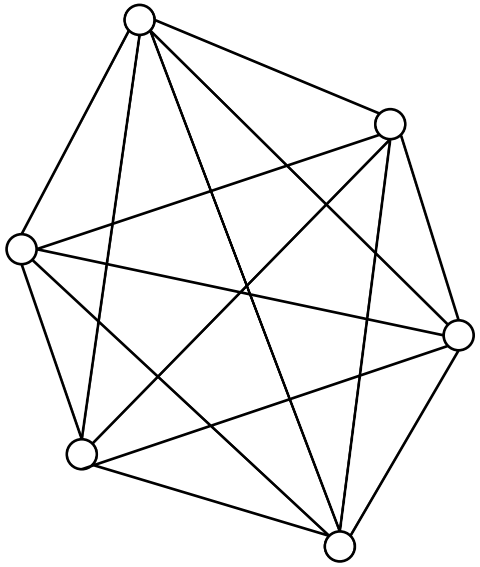

V = states of nature

= strategies of R

= strategies of C

e.g., if $G = K_6$,

θ picked uniformly from vertices

Case I: Signaling scheme reveals θ

- player C picks θ as strategy
- Player R's payoff is 0

Case II: Signaling scheme reveals nothing

- R picks uniformly from vertices
- Player R's payoff is $\geq 1 - 1/3$

Thus, cliques are good for R

Network Security Games

Theorem: Threshold signaling in Network Security Games is NP-hard (reduction from BCBS)

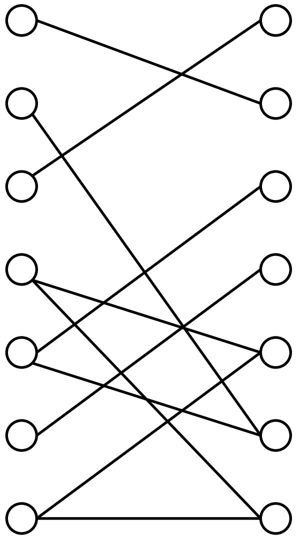

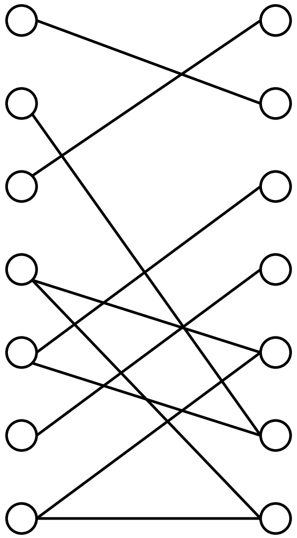

does G contain $K_{r,r}$?

yes, iff

$c = 1 - 1/r^2, \; \exists? \; \mu : c < \mathrm{val}(\mu)$

(in same graph)

Network Security Games

Lemma: G contains $K_{r,r}$ iff there exists μ with $\text{val}(\mu) > c$, $c = 1 - 1/r^2$

(will only sketch one implication)

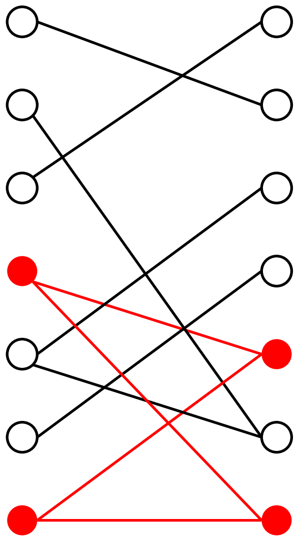

BCBS

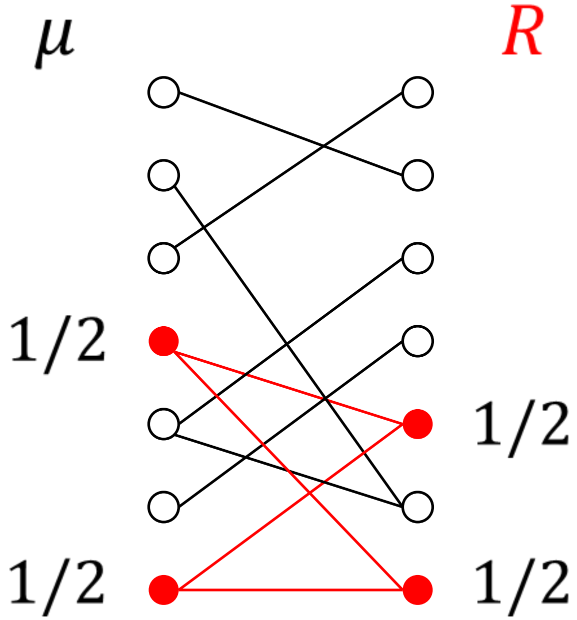

NSG

Say G contains $K_{r,r}$

Choose:

μ uniform distr over one side of $K_{r,r}$

R uniform distr over other side

Then R's payoff is $\geq 1 - 1/r$,

irespective of C's strategy ∎

Network Security Games

Lemma: G contains $K_{r,r}$ iff there exists μ with $\mathrm{val}(\mu) > c$, $c = 1 - 1/r$

+

Theorem: NP-hard to determine if G contains $K_{r,r}$

⇒ Lemma: Threshold Signaling is NP-hard

Network Security Games

Lemma: Threshold Signaling is NP-hard

+

Theorem: Signaling is at least as hard as Threshold Signaling

⇒

Theorem: Signaling in 2-player 0-sum games is NP-hard

Results for 2-player 0-sum games

Result I: NP-hard to obtain an FPTAS

Result II: For a constant $\epsilon > 0$,
computing ϵ-approximate signalling scheme
is as hard as recovering a planted clique in a random graph.

In paper: Signaling in **network congestion games**

Open Questions

Question I: For signaling in 2-player 0-sum games,
ϵ-approximate signalling scheme for some constant $\epsilon > 0$?

Question II: For signaling in 2-player games,
ϵ-approximate signalling scheme for approximate equilibria,
for some constant $\epsilon > 0$?

Question III: What if you could give different signals to different players?
(asymmetric signaling)

Thank You!

Results for 2-player 0-sum games

Result II: For a constant $\epsilon > 0$,
computing ϵ-approximate signalling scheme
is as hard as recovering a planted clique in a random graph.

- why not NP-hard?

$\exists$ quasi-polynomial time algorithm for this [CDDT '15]

so NP-hardness would give a QPT algo for an NP-hard problem

$n^{O(\log n)}$

This Talk

Part I: NP-hard to obtain a *fully polynomial-time approximation scheme*:

for every $\epsilon > 0$, compute ϵ-approximate signalling scheme in time $\text{poly}\left(\frac{1}{\epsilon}\right)$.

Will show:

Dual Linear Program for Signaling

$$\max \sum_{\mu \in \Delta_M} \alpha_\mu \, \mathrm{val}(\mu)$$

$$\text{s.t.} \quad \sum_{\mu \in \Delta_M} \alpha_\mu \, \mu = \lambda$$

$$\alpha_\mu \geq 0$$

Δ_M: set of **all** distributions over states

- this is a linear program
- consider the dual linear program:

$$\min \; w^T \lambda$$

$$w^T \mu \geq \mathrm{val}(\mu) \qquad \text{for all } \mu \in \Delta_M$$

Results for 2-player 0-sum games

Result I: NP-hard to compute **optimal** signaling scheme

NP-hard to obtain a **fully polynomial-time approximation scheme**:

an algorithm that given $\epsilon > 0$,
computes ϵ-approximate signalling scheme
in time poly$\left(\frac{1}{\epsilon}\right)$.

Results for 2-player 0-sum games

Result I: NP-hard to compute **optimal** signaling scheme

For game with m strategies, states of nature,
NP-hard to compute $1/m^8$ - approximate signalling scheme

NP-hard to obtain a FPTAS

an algorithm that given $\epsilon > 0$,
computes ϵ-approximate signalling scheme
in time poly$\left(\frac{1}{\epsilon}\right)$.

This Talk

Part I: For game with m strategies, states of nature,
NP-hard to compute $1/m^8$ - approximate signalling scheme

Part II: For game with m strategies, states of nature,
computing $1/\log^2(m)$-approximate signalling scheme
is as hard as planted-clique recovery

[Dughmi '14]