$\bar{\sigma}_i$ = MAXMIN STRATEGY FOR i $\hat{\sigma}_i$ = MINMAX STRATEGY FOR i

by def. of maxmin we know $u_i(\bar{\sigma}_i, \sigma_{-i}) \geq v_i \quad \forall \sigma_{-i}$ 1)

by def. of minmax we know $u_i(\sigma_i, \hat{\sigma}_{-i}) \leq h_i \quad \forall \sigma_i$ 2)

if 1) then $u_i(\bar{\sigma}_i, \hat{\sigma}_{-i}) \geq v_i$

if 2) then $u_i(\bar{\sigma}_i, \hat{\sigma}_{-i}) \leq h_i$

$\Longrightarrow v_i \leq h_i$

(*A → 43)

to understand why equality holds we need to talk about strong duality

- A NASH EQUILIBRIUM STRATEGY CANNOT BE WORSE THAN A MAXMIN STRATEGY

$\sigma^* = (\sigma_i^*, \sigma_{-i}^*)$ NASH EQ. (we are still in 2-player games)

by def. of NE $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \quad \forall \sigma_i$ (i does not deviate)

$$u_i(\sigma_i, \sigma_{-i}^*) \geq \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) \quad \forall \sigma_i$$

then

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i}) \quad \forall \sigma_i$$

if it holds for any σ_i then it holds also for σ_i that maximizes the right term:

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq \underbrace{\max_{\sigma_i} \min_{\sigma_{-i}} u_i(\sigma_i, \sigma_{-i})}_{\text{maxmin strategy}} = v_i$$

There is a class of games where extremely cautious ① and extremely aggressive ② behaviors spontaneously emerge as rational behaviors.

these are games that encode an interaction setting where assuming that the other just wants to harm IS REAL.

(STRICTLY COMPETITIVE)
ZERO-SUM (TWO-PLAYER) GAMES:

the payoff of the other ≡ my damage!

In these games there's a beautiful relationship between

- Rationality
- Extremely cautious behavior
- Extremely aggressive behavior

MINMAX

MAXMIN

NASH

they become the same thing!

[redacted]

Let's formally analize this statement

$\bar{\sigma}_i$ = maxmin strategy v_i = maxmin value Notation

$\hat{\sigma}_i$ = minmax strategy ($\bar{v}_i$ = minmax value)

σ_i^* = Nash v_i^* = Nash

(A) IF (σ_1^*, σ_2^*) IS NE THEN σ_1^* IS A MAXMIN STRATEGY FOR 1, σ_2^* IS A MAXMIN STRATEGY FOR 2, AND $u_1(\sigma^*) = v_1 = \bar{v}_1$

Def of NE $\quad u_2(\sigma_1^*, \sigma_2^*) \geq u_2(\sigma_1^*, \sigma_2) \qquad \forall \sigma_2$

$$u_2 = -u_1$$

$$u_1(\sigma_1^*, \sigma_2^*) \leq u_1(\sigma_1^*, \sigma_2) \quad \forall \sigma_2$$

so

$$u_1(\sigma_1^*, \sigma_2^*) = \min_{\sigma_2} u_1(\sigma_1^*, \sigma_2)$$

and

$$u_1(\sigma_1^*, \sigma_2^*) \leq \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2)$$

from the previous discussion on the fact that NE $\geq$ maxmin we conclude

$$u_1(\sigma_1^*, \sigma_2^*) = \max_{\sigma_1} \min_{\sigma_2} u_1(\sigma_1, \sigma_2) \longrightarrow \text{EVERY NE YIELDS THE SAME PAYOFF}$$

(B) $(\bar{\sigma}_1, \bar{\sigma}_2)$ is a NASH EQUILIBRIUM

Let's call $V_1 = h_1 = V$

then $V_2 = h_2 = -V$:

why?

Let's see this in SC ~~zero~~ games:

$$V_2 = \max_{\sigma_2} \min_{\sigma_1} M_2(\sigma_1, \sigma_2) \Rightarrow \left(\text{recall that } \max_x f(x) = -\min_x (-f(x)) \right)$$

$$= \max_{\sigma_2} \min_{\sigma_1} \left(-M_1(\sigma_1, \sigma_2)\right) = \max_{\sigma_2} \left(-\max_{\sigma_1} M_1(\sigma_1, \sigma_2)\right)$$

$$= -\min_{\sigma_2} \max_{\sigma_1} M_1(\sigma_1, \sigma_2) = -h_1 = -V_1 = -V$$

$\bar{\sigma}_1$ and $\bar{\sigma}_2$ are maxmin strategies

1) $M_1(\bar{\sigma}_1, \sigma_2) \geq V \quad \forall \sigma_2$ so $M_1(\bar{\sigma}_1, \bar{\sigma}_2) \geq V$

2) $M_2(\sigma_1, \bar{\sigma}_2) \geq -V \quad \forall \sigma_1$ $M_2(\bar{\sigma}_1, \bar{\sigma}_2) \geq -V \rightarrow M_1(\bar{\sigma}_1, \bar{\sigma}_2) \leq V$

$\Rightarrow M_1(\bar{\sigma}_1, \bar{\sigma}_2) = V$ 3)

1), 3) $\Rightarrow M_1(\bar{\sigma}_1, \sigma_2) \geq M_1(\bar{\sigma}_1, \bar{\sigma}_2) \quad \forall \sigma_2$

$\downarrow$

$M_2(\bar{\sigma}_1, \sigma_2) \leq M_2(\bar{\sigma}_1, \bar{\sigma}_2) \quad \forall \sigma_2$ No unilateral deviation of player 2

2), 3) $\Rightarrow M_2(\sigma_1, \bar{\sigma}_2) \geq -M_1(\bar{\sigma}_1, \bar{\sigma}_2) \quad \forall \sigma_1$

$\downarrow$

$M_1(\sigma_1, \bar{\sigma}_2) \leq M_1(\bar{\sigma}_1, \bar{\sigma}_2) \quad \forall \sigma_1$ No unilateral deviation of player 1

NASH

*A why $r_i \leq h_i$, MAXMIN VALUE OF i $\leq$ MINMAX VALUE OF i

- Interpretation based on the notion of commitment:

a player i commits to a strategy σ_i

$\downarrow$

It means: PLAYER i ANNOUNCES TO THE WORLD THAT IT WILL PLAY ACCORDING TO σ_i AND IT WILL ACTUALLY DO IT

PLAYER $-i$ OBSERVES THIS AND EXPLOITS SUCH KNOWLEDGE IN SELECTING THE σ_{-i} [redacted] (Usually a best response)

- ↳ Maxmin assumption is substantially equal to this scenario.
Notice that we are talking about strategies not actions. Actions are observed only when the game is played!

MAXMIN is like:	MINMAX is like:
i commits to a σ	$-i$ commits to a σ_{-i}
$-i$ observes the commitment and selects the σ_{-i} that causes the greatest harm to i	i observes the commitment and selects the σ_i which is a best response

FROM i'S POINT OF VIEW THIS SITUATION IS MORE ADVANTAGEOUS THAN THIS

that's why $r_i \leq h_i$

there's something more actually:

- the full story:

  in two-player games $v_i = h_i$
  (many ways to see this, e.g. duality argument)

  in n-player games $v_i \leq h_i$

  ↳ <u>Team games</u> (we will see later)

Part A) (σ_1^*, σ_2^*) NE $\Rightarrow$ σ_1^* and σ_2^* are maxmin

Part B) $(\bar{\sigma}_1, \bar{\sigma}_2)$ Maxmin $\Rightarrow$ NE

Corollaries:

- any NE in a two-player zero-sum games yields the same payoff
  $\longrightarrow$ WE DO NOT HAVE THE EQUILIBRIUM SELECTION PROBLEM!

- Interchangeability:

(σ_1, σ_2) NE $\qquad (\sigma_1', \sigma_2')$ NE

then by A

σ_1 and σ_1' are maxmin strategies for 1

σ_2 and σ_2' are maxmin strategies for 2

so by B

(σ_1, σ_2') NE, (σ_1', σ_2) NE

We can get the intuition on why it works for minmax as well by thinking about duality ... left for thinking ...

(see (46))

LET'S SEE HOW TO COMPUTE THESE SOLUTION CONCEPT WE DISCUSSED SO FAR.

WE START FROM THE SIMPLEST CASE: ZERO-SUM two player games. Thus, we give a formulation to compute MINMAX = MAXMIN = NE

It's a very simple linear program

we take the stance of player 2, seeking her MINMAX strategy

$$\min \; h_1$$

$$\text{s.t.}$$

$$\sum_{a_2 \in A_2} u_1(a_1, a_2)\, \sigma_2(a_2) \leq h_1 \qquad \forall a_1 \in A_1$$

$$\sum_{a_2 \in A_2} \sigma_2(a_2) = 1$$

$$\sigma_2(a_2) \geq 0 \qquad \forall a_2 \in A_2$$

VARIABLES:
$h_1, \;\; \sigma(a_2) \; \forall a_2 \in A_2$

the utility of 1 should be upper bounded by h_1, no matter what she does

$\downarrow$

I WANT THE SMALLEST OF SUCH BOUNDS!

this is an LP, we can solve it in P time.

We can also write down the dual formulation

It would be a max problem over the lower bound I can guarantee to the player

NOTICE THAT:

- it doesn't matter which player we consider, the formulations are the same. But if we want the complete strategy profile we need to solve one problem for player 1 and one for player 2

- the problem formulation for player i does not contain variables encoding the strategy of player $-i$! IT'S THE "no matter what the other player does" THING WE HAVE IN THE MAXMIN AND MINMAX DEFINITIONS

- Can we apply this formulation to non SC games? YES BUT... WE DON'T GET AN EQUILIBRIUM...

- Can we apply this/these formulations to games with n players?

MAXMIN: YES (easily) BUT...

- no equilibrium
- $v_i < h_i$!

MINMAX: YES but we need a more general definition

$\hat{\sigma}_i^j$:= MinMax strategy of player i AGAINST player j

and...

we need to assume some kind of correlation between other players

→ INTERESTING DIGRESSION: TEAM GAMES

What's a team: A GROUP OF AGENTS WITH THE SAME UTILITY FUNCTION

Careful! They have the same preferences about the world but they are separate players!

What does this mean in game theoretical terms?...

→ NO CORRELATION: THEY CAN PLAN TOGETHER

THEY <u>CANNOT</u> ACT TOGETHER

Special case: STSA-SC

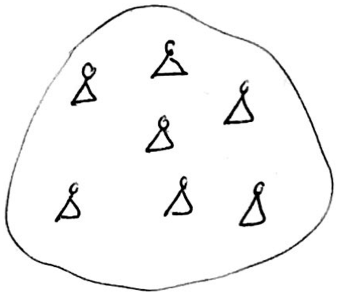

SINGLE TEAM

(n members)

SINGLE ADVERSARY

Zero Sum:

if u_T is the utility of each member of the team then in an outcome x we have

$$(\underbrace{n u_t(x)}_{\text{to the whole team}}, \underbrace{-n u_t(x)}_{\text{to the adversary}})$$

INTERPRETATION:
the adversary "pays" the same amount to each member of the team

- LET'S LOOK AT THE GAME FROM A GENERAL (NOT VERY FORMAL) POINT OF VIEW (see the works of Von Stengel...) WITH MAXMIN IN OUR MINDS.

What does it mean that agents can plan together but then cannot act together. Let's understand this with a simple PE game

TEAM = {Guard 1, Guard 2} ADVERSARY = {Evader}

environment

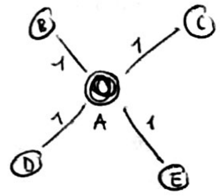

Guards start from A

The evader is somewhere in the graph and
- CAN SEE THE GUARDS AT ANY TIME
- HAS INFINITE SPEED

the team payoff is $\frac{1}{D}$,

D = maximum distance travelled by a guard

- the guards gather in $\overset{A}{\bullet}$ to find a search strategy

[redacted]

Commander: "Let's not risk, we go for a maxmin strategy"

Soldier: "Ok, I have a strategy that guarantees us $\geq \frac{1}{3}$"

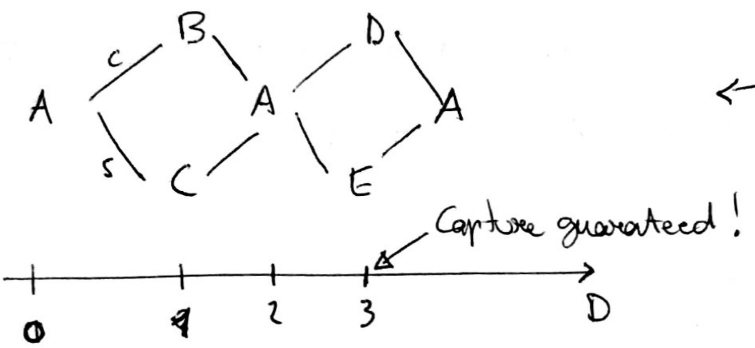

← Strategy S^*

Commander: [redacted]

: "We need to operate in stealth mode"
(invisible to each other, radio silence, etc...)

Soldier: "Ok, so we [redacted] must settle for $\geq \frac{1}{7}$"

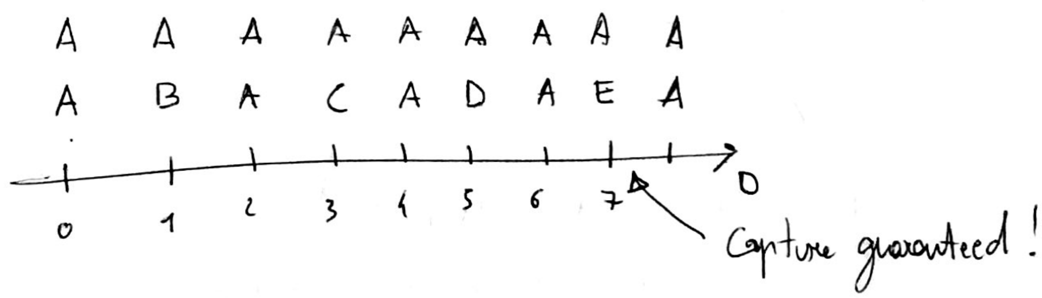

→ Careful! Modeling this with a strategic game is not easy but can be done.

Still, we can see how maxmin and minmax rationales apply here:

$\frac{1}{3}$ is the maxmin as if the team is a single player

$\frac{1}{7}$ is the maxmin as if the team are two separate players, but with the same objectives and aware of each other

$\Longrightarrow$ THE NON-AVAILABILITY OF CORRELATION GIVES US A WORSE GUARANTEE ON THE MINIMUM UTILITY WE CAN ACHIEVE!

GETTING $\geq \frac{1}{3}$ WILL ENTAIL COORDINATION "OUTSIDE THE STRUCTURE OF THE GAME" (It would be another game! Not this one)

What's happening on the other side??

Evader: "I want them to travel the largest distance, I don't care if they catch me at some point"

↳ IT'S A MINMAX PLAYER (Let's assume it...)

"Given that their best case is S^*, I will play such that $\leq \frac{1}{3}$"

(the minmax player BY DEFINITION does not account for the absence of correlation)

BUT

$$\frac{1}{7} \neq \frac{1}{3}$$

It's the team maxmin

IT'S NOT AN EQUILIBRIUM!! (despite being a SC competitive setting)

But there are some good news (and a bad one)

- NASH: an equilibrium always exists (we know this)

- VON STENGEL and KOLLER (1997)

  - In at least one NE the team plays a team-maxmin strategy, let's call it TEAM-MAXMIN-NE

How to compute it?

$$\left[\begin{array}{l} \text{TEAM} = \{1, 2, \dots, n\} = T \\ n+1 \text{ is the adversary} \end{array}\right]$$

$$\max \ v$$

$$v \leq \sum_{a_T \in A_T} u_t(a_T, a_{n+1}) \prod_{i \in T} \sigma_i(a_i) \qquad \forall a_{n+1} \in A_{n+1}$$

non correlation

$$\sum_{a_i \in A_i} \sigma_i(a_i) = 1 \qquad \forall i \in T \cup \{n+1\}$$

$$\sigma_i(a_i) \geq 0 \quad \forall i \in T \cup \{n+1\},\ a_i \in A_i$$

It's a non linear problem .... :(

  - No other NE can provide a better utility to the team (good! We don't have the equilibrium selection problem!)

- OUR (VERY) RECENT RESULT: It can be approximated within any additive error $\epsilon \geq \sqrt{\frac{\ln m}{2n}}$ in quasi-polinomial time

$$\left(m^{O\left(\frac{(n)(\ln m)}{2\epsilon^2}\right)}\right)$$

- EMPIRICALLY: BETTER GO WITH GLOBAL OPTIMIZATION

LET'S NOW TURN TO NASH EQUILIBRIUM IN GENERAL SM GAMES (TWO PLAYERS)

We'll see 3 methods or, more properly, approaches

- LEMKE-HOWSON (LH): linear complementarity
- PORTER-NUDELMAN-SHOAM (PNS): support enumeration
- SANDHOLM-GILPIN-CONITZER (SGC): MIP-NASH

• A note on complexity

- Discussing the worst case computational complexity can be quite technical because of Nash theorem

- VERY VERY SHORT STORY: "traditional" complexity analysis makes use of decision problem; ACTUALLY, THE WHOLE THEORY OF NP-COMPLETENESS ONLY APPLIES TO DECISION PROBLEM

Decision problem: YES/NO PROBLEMS DERIVED FROM OPTIMIZATION COUNTERPARTS (the rationale being: the decision problem is not harder than its optimization counterpart)

NE is a "always-yes" problem!

- No NP-Completess reduction discovered to date in the general case, some NP-Hardness results for particular cases...
- Complete wrt another class of problems (which share the feature of always having a "yes" answer): PPAD "polynomial parity argument, directed version"

  ⟶ Megiddo, Papadimitriou, Goldberg, Daskalakis 2006!

generally believed that PPAD $\neq$ P

• Lemke – Howson (Linear Complementarity Formulation)

U_i^* = utility that player i gets at the equilibrium

FOR PLAYER i, THE EXPECTED UTILITY OF PLAYING ACTION a_i CAN ALWAYS BE EXPRESSED AS $U_i^* - x^i$ — a slack quantity that depends on a_i. ALSO CALLED "REGRET"

Formulation LH-1

$$\sum_{a_2 \in A_2} u_1(a_j, a_2)\, \underbrace{\sigma_2(a_2)}_{\text{variables}} = \underbrace{U_1^*}_{\text{variables}} - \underbrace{x_1^j}_{\text{variables}} \qquad \forall\, a_j \in A_1$$

$$\sum_{a_1 \in A_1} u_2(a_1, a_j)\, \sigma_1(a_1) = U_2^* - x_2^j \qquad \forall\, a_j \in A_2$$

$$\sigma_i(a_i) \geq 0 \quad \forall i,\, a_i \in A_i \qquad \sum_{a_i \in A_i} \sigma_i(a_i) = 1 \quad \forall i$$

$$x_i^j \geq 0 \quad \forall i,\, a_j \in A_i$$

NOW RECALL THAT AT THE EQUILIBRIUM THE AGENTS MUST BE INDIFFERENT OVER THE ACTIONS THEY PLAY WITH POSITIVE PROBABILITY

In our formulation:

if $\sigma_i(a_j) > 0$ then $x_i^j = 0$

if $\sigma_i(a_j) = 0$ then $x_i^j \geq 0$

In other words x_i^j is the incentive of player i of deviating from action a_j

LH-2

$$x_i^j \cdot \sigma_i(a_j) = 0 \qquad \forall i,\, a_j \in A_j$$

non linear constraint (linear complementarity) the best algorithm to solve an LC problem is Lemke–Howson

Unfortunately, there's no time to see it but here are some issues worth noting:

- It runs in exponential time (worst case)
- the way it solves the problem can give deep algorithmic insights on the (algorithmic) nature of Nash (so if you have time, give it a look!) $\rightsquigarrow$ NE $\underline{\underline{\text{IS}}}$ A COMPLEMENTARITY PROBLEM

• PORTER - NUDELMAN - SHOHAM (support enumeration)

Heuristic approach based on the following rationale:

GIVEN A SUPPORT PROFILE $S = (S_1, S_2)$

(where $S_i \subseteq A_i$)

TESTING WETHER THERE EXIST A NE WITH SUPPORT S CAN BE DONE BY A SIMPLE LINEAR (FEASIBILITY PROGRAM). THIS ONE

$$\sum_{a_{-i} \in S_{-i}} \sigma_{-i}(a_{-i}) \, u_i(a_i, a_{-i}) = v_i \qquad \forall i \in \{1,2\}, \ \underbrace{a_i \in S_i}$$

actions within the support: THE PLAYER MUST BE INDIFFEREN

$$\sum_{a_{-i} \in S_{-i}} \sigma_{-i}(a_{-i}) \, u_i(a_i, a_{-i}) \leq v_i \qquad \forall i \in \{1,2\}, \ \underbrace{a_i \notin S_i}$$

actions outside the support: THEIR UTILITY CANNOT BE GREATER THAN THOSE IN THE SUPPORT then

$$\sigma_i(a_i) \geq 0 \quad \forall i \in \{1,2\}, \; a_i \in S_i$$

$$\sigma_i(a_i) = 0 \quad \forall i \in \{1,2\}, \; a_i \notin S_i$$

actions outside the support must be played with prob. zero (by definition of support)

$$\sum_{a_i \in A_i} \sigma_i(a_i) = 1 \quad \forall i \in \{1,2\}$$

Simple algorithm: ENUMERATE ALL SUPPORT PROFILES AND, FOR EACH OF THOSE, CHECK THE ABOVE FEASIBILITY PROGRAM

→ The supports profile are exponentially many!

Heuristic principles

- FAVOUR BALANCED SUPPORT PROFILES AGAINST NON BALANCED ONES (each player has the same # of actions)
- FAVOUR SMALL SUPPORTS OVER BIG ONES
- USE DOMINANCE (STRICT) FOR PRUNING

→ Empirically PNS outperforms LH

- Sandholm – Gilpin – Conitzer

We use the same idea from LH but instead of the linear complementarity constraint we use binary variables

In detail, we are going to introduce

$$b_i^j = \begin{cases} 1 & a_j \in A_i \text{ is NOT in the support} \\ 0 & \text{otherwise} \end{cases}$$

M a large constant

We rewrite formulation LH-1 with the following additional constraints (replacing linear complementarity)

a) $\sigma_i(a_j) \leq 1 - b_i^j \quad \forall i \in \{1,2\}, a_j \in A_i$

b) $r_i^j \leq M b_i^j \quad \forall i \in \{1,2\}, a_j \in A_i$

from a we get the definition of b: $b_i^j = 1 \iff \sigma_i(a_j) = 0$

from b we get what we previously had from linear complementarity:

if $b_i^j = 1$ the regret can be unbounded

else it MUST be zero!

We get rid of the variable multiplication but we have the integer variable (NP-Hard problem) there are variants of this approach that make use of an objective function

- we can embed feasibility or even equilibrium selection

what is the best method? The answer is clearly not possible, but if I have to pick one ... LH.

$\hat{\sigma}_1, \hat{\sigma}_2$ are minmax strategies $\Rightarrow$ NE

1) $u_1(\sigma_1, \hat{\sigma}_2) \leq v \quad \forall \sigma_1 \quad \Rightarrow u_1(\hat{\sigma}_1, \hat{\sigma}_2) \leq v$

2) $u_2(\hat{\sigma}_1, \sigma_2) \leq -v \quad \forall \sigma_2 \quad \Rightarrow u_2(\hat{\sigma}_1, \hat{\sigma}_2) \leq -v \Rightarrow u_1(\hat{\sigma}_1, \hat{\sigma}_2) \geq v$

3) $\Rightarrow u_1(\hat{\sigma}_1, \hat{\sigma}_2) = v$

1) 3) :

$u_1(\sigma_1, \hat{\sigma}_2) \leq u_1(\hat{\sigma}_1, \hat{\sigma}_2) \quad \forall \sigma_1$ no unilateral deviation of 1

2) 3) :

$u_2(\hat{\sigma}_1, \sigma_2) \leq -u_1(\hat{\sigma}_1, \hat{\sigma}_2) \quad \forall \sigma_2$

$u_2(\hat{\sigma}_1, \sigma_2) \leq u_2(\hat{\sigma}_1, \hat{\sigma}_2) \quad \forall \sigma_2$ no unilateral deviation of 2

Corollary of

(x) player 1 :

If she plays $\bar{\sigma}_1$ then $u_1(\bar{\sigma}_1, \sigma_2) \geq v \quad \forall \sigma_2$

If she plays $\hat{\sigma}_1$ then $u_2(\hat{\sigma}_1, \sigma_2) \leq -v \quad \forall \sigma_2$

but then

$u_1(\hat{\sigma}_1, \sigma_2) \geq v \; \forall \sigma_2$ so $\hat{\sigma}_1 = \bar{\sigma}_1$
