Epistemic Game Theory

Lecture 4

Eric Pacuit

University of Maryland, College Park
pacuit.org
epacuit@umd.edu

February 17, 2014

Let $G = \langle \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a finite strategic game (each S_i is finite and the set of players N is finite).

A **strategy profile** is an element $\sigma \in S = S_1 \times \cdots \times S_n$

σ is a **Nash equilibrium** provided for all i, for all $s_i \in S_i$,

$$u_i(\sigma) \geq u_i(s_i, \sigma_{-i})$$

Let $G = \langle \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a finite strategic game.

$\Sigma_i = \{p \mid p : S_i \to [0,1] \text{ and } \sum_{s_i \in S_i} p(s_i) = 1\}$

The **mixed extension** of G is the game $\langle \{\Sigma_i\}_{i \in N}, \{U_i\}_{i \in N} \rangle$ where for $\sigma \in \Sigma = \Sigma_1 \times \cdots \times \Sigma_n$:

$$U_i(\sigma) = \sum_{(s_1,\ldots,s_n) \in S} \sigma_1(s_1)\sigma_2(s_2)\cdots\sigma_n(s_n)u_i(s_1,\ldots,s_n)$$

Theorem. Suppose that σ is a Nash equilibrium in mixed strategies for a game $G = \langle \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$. Suppose that $s_i, s_i^* \in S_i$ are two pure strategies such that $\sigma_i(s_i) > 0$ and $\sigma_i(s_i^*) > 0$, then

$$U_i(s_i, \sigma_{-i}) = U_i(s_i^*, \sigma_{-i})$$

Theorem (Nash). Every finite game G has a Nash equilibrium in mixed strategies (i.e., there is a Nash equilibrium in the mixed extension G).

Not all equilibrium are created equal...

Perfect equilibrium

		Bob	
		L	R
Ann	U	1,1	0,0
	D	0,0	0,0

Perfect equilibrium

		Bob L	Bob R
Ann	U	1,1	0,0
Ann	D	0,0	0,0

Perfect equilibrium

		Bob	
		L	R
Ann	U	1,1	0,0
	D	0,0	0,0

Isn't (U, L) more "reasonable" than (D, R)?

Perfect equilibrium

		Bob	
		L	R
Ann	U	1,1	0,0
	D	0,0	0,0

Completely mixed strategy: a mixed strategy in which every strategy gets some positive probability

Perfect equilibrium

		Bob	
		L	R
Ann	U	1,1	0,0
	D	0,0	0,0

Completely mixed strategy: a mixed strategy in which every strategy gets some positive probability

ϵ**-perfect equilibrium**: a completely mixed strategy profile in which any pure strategy that is not a best reply receives probability less than ϵ

Prefect equilibrium: the mixed strategy profile that is the limit as ϵ goes to 0 of ϵ-prefect equilibria.

Proper equilibrium

		Bob L	Bob C	Bob R
Ann	U	-9,-9	-7,-7	-7,-7
Ann	M	0,0	0,0	-7,-7
Ann	D	1,1	0,0	-9,-9

Proper equilibrium

		Bob	
	L	C	R
U	-9,-9	-7,-7	-7,-7
Ann M	0,0	0,0	-7,-7
D	1,1	0,0	-9,-9

Proper equilibrium

		Bob L	C	R
Ann	U	-9,-9	-7,-7	-7,-7
	M	0,0	0,0	-7,-7
	D	1,1	0,0	-9,-9

ϵ-**proper equilibrium**: a completely mixed strategy profile such that if strategy s is a better response than s', then $\frac{p(s)}{p(s')} < \epsilon$

Proper equilibrium: the mixed strategy profile that is the limit as ϵ goes to 0 of ϵ-proper equilibria.

Proper equilibrium

		Bob		
		L	C	R
Ann	U	-9,-9	-7,-7	-7,-7
	M	0,0	0,0	-7,-7
	D	1,1	0,0	-9,-9

ϵ-**proper equilibrium**: a completely mixed strategy profile such that if strategy s is a better response than s', then $\frac{p(s)}{p(s')} < \epsilon$

Proper equilibrium: the mixed strategy profile that is the limit as ϵ goes to 0 of ϵ-proper equilibria.

Normal form vs. Extensive form

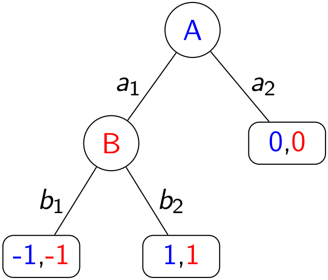

Normal form vs. Extensive form

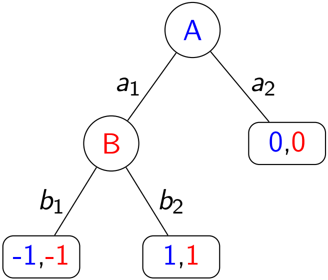

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

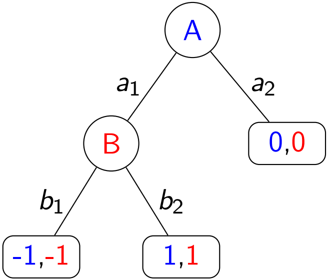

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

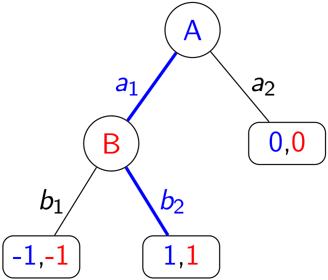

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

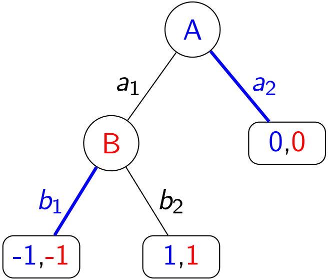

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

(Cf. the various notions of *sequential equilibrium*)

T. Seidenfeld. *When normal and extensive form decisions differer*. in *Logic, Methodology and Philosophy of Science IX*, Elsevier, 1994.

Trembling Hands

"There cannot be any mistakes if the players are absolutely rational. Nevertheless, a satisfactory interpretation of equilibrium points in extensive games seems to require that the possibility of mistakes is not completely excluded. This can be achieved by a point of view which looks at complete rationality as the limiting case of incomplete rationality." (pg. 35)

R. Selten. *Reexamination of the Perfectness Concept of Equilibrium in Extensive Games*. International Journal of Game Theory, 4, pgs. 25 - 55, 1975.

Why play Nash equilibrium?

Self-Enforcing Agreements: Nash equilibria are recommended by being the only strategy combinations on which the players could make self-enforcing agreements, i.e., agreements that each has reason to respect, even without external enforcement mechanisms.

M. Risse. *What is rational about Nash equilibria?*. Synthese, 124:3, pgs. 361 - 384, 2000.

Some equilibria are not self-enforcing

Bob

		L	R
Ann	U	9,9	0,8
	D	8,0	7,7

Some equilibria are not self-enforcing

		Bob L	Bob R
Ann	U	9,9	0,8
Ann	D	8,0	7,7

Some equilibria are not self-enforcing

		Bob	
		L	R
Ann	U	9,9	0,8
	D	8,0	7,7

An agreement to play (U, U) is *not* self-enforcing: Ann has a good reason to believe that Bob will deviate (similarly for Bob)

Some non-equilibria are self-enforcing

Bob

		L	R
Ann	U	0,0	4,2
	D	2,4	3,3

Some non-equilibria are self-enforcing

		Bob	
		L	R
Ann	U	0,0	4,2
	D	2,4	3,3

Some non-equilibria are self-enforcing

Bob

Ann		L	R
	U	0,0	4,2
	D	2,4	3,3

An agreement to play (D, D) is self-enforcing: Both risk ending up with 0 if they deviate from the agreement.

Towards an Epistemic Characterization of Nash Equilibria

Correlation: Players can improve their expected value by correlating their choices on an "outside signal"

Correlated Strategies

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

- Three Nash equilibria:
  - (U, L): the payoff is $(2, 1)$
  - (D, R): the payoff is $(1, 2)$
  - $([\frac{2}{3}(U), \frac{1}{3}D], [\frac{1}{3}(L), \frac{2}{3}(R)])$: the payoff is $(\frac{2}{3}, \frac{2}{3})$

Correlated Strategies

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

- Three Nash equilibria:
  - (U, L): the payoff is $(2, 1)$
  - (D, R): the payoff is $(1, 2)$
  - $([\frac{2}{3}(U), \frac{1}{3}D], [\frac{1}{3}(L), \frac{2}{3}(R)])$: the payoff is $(\frac{2}{3}, \frac{2}{3})$
- Mixed Strategies: Each player conducts a private, independent lottery to choose their strategy.

Correlated Strategies

	L	R
U	2, 1	0, 0
D	0, 0	1, 2

	L	R
U	0.5	0
D	0	0.5

- Three Nash equilibria:
  - (U, L): the payoff is $(2, 1)$
  - (D, R): the payoff is $(1, 2)$
  - $([\frac{2}{3}(U), \frac{1}{3}D], [\frac{1}{3}(L), \frac{2}{3}(R)])$: the payoff is $(\frac{2}{3}, \frac{2}{3})$
- Mixed Strategies: Each player conducts a private, independent lottery to choose their strategy.
- Conduct a *public* lottery: flip a fair coin and follow the strategy $(H \Rightarrow (U, L),\ T \Rightarrow (D, R))$. The payoff is $(1.5, 1.5)$.

Two extremes:

1. Completely private, independent lotteries
2. A single, completely public lottery

Two extremes:

1. Completely private, independent lotteries
2. A single, completely public lottery

What about: a public lottery, but reveal only partial information about the outcome to each of the players?

A	L	R
U	0, 1, 3	0, 0, 0
D	1, 1, 1	1, 0, 0

B	L	R
U	2, 2, 2	0, 0, 0
D	2, 2, 0	2, 2, 2

C	L	R
U	0, 1, 0	0, 0, 0
D	1, 1, 1	1, 0, 3

- Three player game: Ann chooses the row, Bob chooses the column, Charles chooses the matrix

A	L	R
U	0, 1, 3	0, 0, 0
D	1, 1, 1	1, 0, 0

B	L	R
U	2, 2, 2	0, 0, 0
D	2, 2, 0	2, 2, 2

C	L	R
U	0, 1, 0	0, 0, 0
D	1, 1, 1	1, 0, 3

- Three player game: Ann chooses the row, Bob chooses the column, Charles chooses the matrix

A	L	R
U	0, 1, 3	0, 0, 0
D	1, 1, 1	1, 0, 0

B	L	R
U	2, 2, 2	0, 0, 0
D	2, 2, 0	2, 2, 2

C	L	R
U	0, 1, 0	0, 0, 0
	1, 1, 1	1, 0, 3

- Three player game: Ann chooses the row, Bob chooses the column, Charles chooses the matrix

A	L	R
U	0, 1, 3	0, 0, 0
D	1, 1, 1	1, 0, 0

B	L	R
U	2, 2, 2	0, 0, 0
D	2, 2, 0	2, 2, 2

C	L	R
U	0, 1, 0	0, 0, 0
	1, 1, 1	1, 0, 3

- Three player game: Ann chooses the row, Bob chooses the column, Charles chooses the matrix

A	L	R
U	0, 1, 3	0, 0, 0
D	1, 1, 1	1, 0, 0

B	L	R
U	2, 2, 2	0, 0, 0
D	2, 2, 0	2, 2, 2

C	L	R
U	0, 1, 0	0, 0, 0
D	1, 1, 1	1, 0, 3

- Three player game: Ann chooses the row, Bob chooses the column, Charles chooses the matrix
- The only equilibrium payoff is $(1, 1, 1)$

A	L	R
U	0, 1, 3	0, 0, 0
D	1, 1, 1	1, 0, 0

B	L	R
U	2, 2, 2	0, 0, 0
D	2, 2, 0	2, 2, 2

C	L	R
U	0, 1, 0	0, 0, 0
D	1, 1, 1	1, 0, 3

- Three player game: Ann chooses the row, Bob chooses the column, Charles chooses the matrix
- The only equilibrium payoff is $(1, 1, 1)$
- There is a correlated mechanism that produces $(2, 2, 2)$
  - Ann and Bob toss a fair coin, but do not reveal the result to Charles
  - Ann and Bob correlate their choices on the coin toss $(H \Rightarrow (U, L), T \Rightarrow (D, R))$
  - Charles choose B

	L	R
U	6, 6	2, 7
D	7, 2	0, 0

- Three Nash equilibria:
  - (U, R): the payoff is $(2, 7)$
  - (D, L): the payoff is $(7, 2)$
  - $([\frac{2}{3}(U), \frac{1}{3}D], [\frac{2}{3}(L), \frac{1}{3}(R)])$: the payoff is $(4\frac{2}{3}, 4\frac{2}{3})$

	L	R
U	6, 6	2, 7
D	7, 2	0, 0

	L	R
U	$1/3$	$1/3$
D	$1/3$	0

- Three Nash equilibria:
  - (U, R): the payoff is $(2, 7)$
  - (D, L): the payoff is $(7, 2)$
  - $([\frac{2}{3}(U), \frac{1}{3}D], [\frac{2}{3}(L), \frac{1}{3}(R)])$: the payoff is $(4\frac{2}{3}, 4\frac{2}{3})$
- After conducting the lottery, an outside observer provides Ann with a recommendation to play the first component of the profile that was chosen, and Bob the second component.

	L	R
U	6, 6	2, 7
D	7, 2	0, 0

	L	R
U	1/3	1/3
D	1/3	0

- Three Nash equilibria:
  - (U, R): the payoff is $(2, 7)$
  - (D, L): the payoff is $(7, 2)$
  - $([\frac{2}{3}(U), \frac{1}{3}D], [\frac{2}{3}(L), \frac{1}{3}(R)])$: the payoff is $(4\frac{2}{3}, 4\frac{2}{3})$
- After conducting the lottery, an outside observer provides Ann with a recommendation to play the first component of the profile that was chosen, and Bob the second component.
- The expected payoff is $\frac{1}{3}(6, 6) + \frac{1}{3}(2, 7) + \frac{1}{3}(7, 2) = (5, 5)$ (which is outside the convex hull of the Nash equilibria)

Towards an Epistemic Characterization of Nash Equilibria

Correlation: Players can improve their expected value by correlating their choices on an "outside signal"

Towards an Epistemic Characterization of Nash Equilibria

Correlation: Players can improve their expected value by correlating their choices on an "outside signal"

With more than 2 players...

- A player may believe that (some of) the other players strategy choices are **independent** or **correlated**.
- Two players can **agree** or **disagree** on the probabilities that the assign to a third player's choice of strategy.

Comparing Dominance Reasoning and MEU

$G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$
$X \subseteq S_{-i}$ (a set of strategy profiles for all players except i)

Comparing Dominance Reasoning and MEU

$G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$
$X \subseteq S_{-i}$ (a set of strategy profiles for all players except i)

$s, s' \in S_i$, s **strictly dominates** s' with respect to X provided

$$\forall s_{-i} \in X, \quad u_i(s, s_{-i}) > u_i(s', s_{-i})$$

Comparing Dominance Reasoning and MEU

$G = \langle N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$
$X \subseteq S_{-i}$ (a set of strategy profiles for all players except i)

$s, s' \in S_i$, s **strictly dominates** s' with respect to X provided

$$\forall s_{-i} \in X, \quad u_i(s, s_{-i}) > u_i(s', s_{-i})$$

$s, s' \in S_i$, s **weakly dominates** s' with respect to X provided

$$\forall s_{-i} \in X, \ \ u_i(s, s_{-i}) \geq u_i(s', s_{-i}) \ \text{ and } \ \exists s_{-i} \in X, \ \ u_i(s, s_{-i}) > u_i(s', s_{-i})$$

Comparing Dominance Reasoning and MEU

$G = \langle N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$
$X \subseteq S_{-i}$ (a set of strategy profiles for all players except i)

$s, s' \in S_i$, s **strictly dominates** s' with respect to X provided

$$\forall s_{-i} \in X, \quad u_i(s, s_{-i}) > u_i(s', s_{-i})$$

$s, s' \in S_i$, s **weakly dominates** s' with respect to X provided

$$\forall s_{-i} \in X, \; u_i(s, s_{-i}) \geq u_i(s', s_{-i}) \text{ and } \exists s_{-i} \in X, \; u_i(s, s_{-i}) > u_i(s', s_{-i})$$

$p \in \Delta(X)$, s is a **best response** to p with respect to X provided

$$\forall s' \in S_i, \quad EU(s, p) \geq EU(s', p)$$

Strict Dominance and MEU

Fact. Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a strategic game and $X \subseteq S_{-i}$. A strategy $s_i \in S_i$ is strictly dominated (possibly by a mixed strategy) with respect to X iff there is no probability measure $p \in \Delta(X)$ such that s_i is a best response to p.

Suppose that $G = \langle N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$ is a finite strategic game.

Suppose that $G = \langle N, \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$ is a finite strategic game.

Suppose that $s_i \in S_i$ is strictly dominated with respect to X:

$$\exists s_i' \in S_i, \forall s_{-i} \in X, \quad u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$$

Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a finite strategic game.

Suppose that $s_i \in S_i$ is strictly dominated with respect to X:

$$\exists s_i' \in S_i, \forall s_{-i} \in X, \quad u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$$

Let $p \in \Delta(X)$ be any probability measure. Then,

$$\forall s_{-i} \in X, \quad p(s_{-i}) \cdot u_i(s_i', s_{-i}) \geq p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

$$\exists s_{-i} \in X, \quad p(s_{-i}) \cdot u_i(s_i', s_{-i}) > p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a finite strategic game.

Suppose that $s_i \in S_i$ is strictly dominated with respect to X:

$$\exists s_i' \in S_i, \forall s_{-i} \in X, \quad u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$$

Let $p \in \Delta(X)$ be any probability measure. Then,

$$\forall s_{-i} \in X, \quad p(s_{-i}) \cdot u_i(s_i', s_{-i}) \geq p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

$$\exists s_{-i} \in X, \quad p(s_{-i}) \cdot u_i(s_i', s_{-i}) > p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

Hence,

$$\sum_{s_{-i} \in S_{-i}} p(s_{-i}) \cdot u_i(s_i', s_{-i}) > \sum_{s_{-i} \in S_{-i}} p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a finite strategic game.

Suppose that $s_i \in S_i$ is strictly dominated with respect to X:

$$\exists s_i' \in S_i, \forall s_{-i} \in X, \quad u_i(s_i', s_{-i}) > u_i(s_i, s_{-i})$$

Let $p \in \Delta(X)$ be any probability measure. Then,

$$\forall s_{-i} \in X, \quad p(s_{-i}) \cdot u_i(s_i', s_{-i}) \geq p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

$$\exists s_{-i} \in X, \quad p(s_{-i}) \cdot u_i(s_i', s_{-i}) > p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

Hence,

$$\sum_{s_{-i} \in S_{-i}} p(s_{-i}) \cdot u_i(s_i', s_{-i}) > \sum_{s_{-i} \in S_{-i}} p(s_{-i}) \cdot u_i(s_i, s_{-i})$$

So, $EU(s_i', p) > EU(s_i, p)$: s_i is not a best response to p.

For the converse direction, we sketch the proof for two player games and where $X = S_{-i}$. [1]

[1] The proof of the more general statement uses the *supporting hyperplane theorem* from convex analysis.

For the converse direction, we sketch the proof for two player games and where $X = S_{-i}$. [1]

Let $G = \langle S_1, S_2, u_1, u_2 \rangle$ be a two-player game.
(Let $U_i : \Delta(S_1) \times \Delta(S_2) \to \mathbb{R}$ be the expected utility for i)

[1] The proof of the more general statement uses the *supporting hyperplane theorem* from convex analysis.

For the converse direction, we sketch the proof for two player games and where $X = S_{-i}$. [1]

Let $G = \langle S_1, S_2, u_1, u_2 \rangle$ be a two-player game.
(Let $U_i : \Delta(S_1) \times \Delta(S_2) \to \mathbb{R}$ be the expected utility for i)

Suppose that $\alpha \in \Delta(S_1)$ is not a best response to any $p \in \Delta(S_2)$.

$$\forall p \in \Delta(S_2) \;\; \exists q \in \Delta(S_1), \quad U_1(q, p) > U_1(\alpha, p)$$

[1] The proof of the more general statement uses the *supporting hyperplane theorem* from convex analysis.

For the converse direction, we sketch the proof for two player games and where $X = S_{-i}$. [1]

Let $G = \langle S_1, S_2, u_1, u_2 \rangle$ be a two-player game.
(Let $U_i : \Delta(S_1) \times \Delta(S_2) \to \mathbb{R}$ be the expected utility for i)

Suppose that $\alpha \in \Delta(S_1)$ is not a best response to any $p \in \Delta(S_2)$.

$$\forall p \in \Delta(S_2) \;\; \exists q \in \Delta(S_1), \quad U_1(q, p) > U_1(\alpha, p)$$

We can define a function $b : \Delta(S_2) \to \Delta(S_1)$ where, for each $p \in \Delta(S_2)$, $U_1(b(p), p) > U_1(\alpha, p)$.

[1] The proof of the more general statement uses the *supporting hyperplane theorem* from convex analysis.

Consider the game $G' = \langle S_1, S_2, \overline{u}_1, \overline{u}_2 \rangle$ where

$$\overline{u}_1(s_1, s_2) = u_1(s_1, s_2) - U_1(\alpha, s_2) \text{ and } \overline{u}_2(s_1, s_2) = -\overline{u}_1(s_1, s_2)$$

Consider the game $G' = \langle S_1, S_2, \overline{u}_1, \overline{u}_2 \rangle$ where

$$\overline{u}_1(s_1, s_2) = u_1(s_1, s_2) - U_1(\alpha, s_2) \text{ and } \overline{u}_2(s_1, s_2) = -\overline{u}_1(s_1, s_2)$$

By the minimax theorem, there is a Nash equilibrium (p_1^*, p_2^*) such that for all $m \in \Delta(S_2)$,

$$\overline{U}(p_1^*, m) \geq \overline{U}_1(p_1^*, p_2^*) \geq \overline{U}_1(b(p_2^*), p_2^*)$$

Consider the game $G' = \langle S_1, S_2, \overline{u}_1, \overline{u}_2 \rangle$ where

$$\overline{u}_1(s_1, s_2) = u_1(s_1, s_2) - U_1(\alpha, s_2) \text{ and } \overline{u}_2(s_1, s_2) = -\overline{u}_1(s_1, s_2)$$

By the minimax theorem, there is a Nash equilibrium (p_1^*, p_2^*) such that for all $m \in \Delta(S_2)$,

$$\overline{U}(p_1^*, m) \geq \overline{U}_1(p_1^*, p_2^*) \geq \overline{U}_1(b(p_2^*), p_2^*)$$

We now prove that $\overline{U}_1(b(p_2^*), p_2^*) > 0$:

$$\overline{U}_1(b(p_2^*), p_2^*) \quad = \quad \sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) \overline{u}_1(x, y)$$

$$
\begin{aligned}
\overline{U}_1(b(p_2^*), p_2^*) &= \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)\overline{u}_1(x,y) \\
&= \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)[u_1(x,y) - U_1(\alpha, y)]
\end{aligned}
$$

$$
\begin{aligned}
\overline{U}_1(b(p_2^*), p_2^*) &= \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)\overline{u}_1(x,y) \\
&= \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)[u_1(x,y) - U_1(\alpha, y)] \\
&= \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)u_1(x,y) \\
&- \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha, y)
\end{aligned}
$$

$$
\begin{array}{rcl}
\overline{U}_1(b(p_2^*), p_2^*) & = & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)\overline{u}_1(x,y) \\[2ex]
& = & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)[u_1(x,y) - U_1(\alpha, y)] \\[2ex]
& = & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)u_1(x,y) \\
& - & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha, y) \\[2ex]
& = & U_1(b(p_2^*), p_2^*) \\
& - & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha, y)
\end{array}
$$

$$
\begin{aligned}
\overline{U}_1(b(p_2^*), p_2^*) &= \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)\overline{u}_1(x,y) \\
&= \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)[u_1(x,y) - U_1(\alpha,y)] \\
&= \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)u_1(x,y) \\
&- \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y) \\
&= {\color{blue}U_1(b(p_2^*), p_2^*)} \\
&- \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y) \\
&> {\color{blue}U_1(\alpha, p_2^*)} - \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y)
\end{aligned}
$$

$$\begin{aligned}
\overline{U}_1(b(p_2^*), p_2^*) &= \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)\overline{u}_1(x,y) \\
&= \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)[u_1(x,y) - U_1(\alpha,y)] \\
&= \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)u_1(x,y) \\
&- \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y) \\
&= U_1(b(p_2^*), p_2^*) \\
&- \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y) \\
&> U_1(\alpha, p_2^*) - \textstyle\sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y) \\
&= U_1(\alpha, p_2^*) - \textstyle\sum_{x\in S_1} b(p_2^*)(x)\sum_{y\in S_2} p_2^*(y)U_1(\alpha,y)
\end{aligned}$$

$$
\begin{array}{rcl}
\overline{U}_1(b(p_2^*), p_2^*) & = & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)\overline{u}_1(x,y) \\
 & = & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)[u_1(x,y) - U_1(\alpha,y)] \\
 & = & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)u_1(x,y) \\
 & - & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y) \\
 & = & U_1(b(p_2^*), p_2^*) \\
 & - & \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y) \\
 & > & U_1(\alpha, p_2^*) - \sum_{x\in S_1}\sum_{y\in S_2} b(p_2^*)(x)p_2^*(y)U_1(\alpha,y) \\
 & = & U_1(\alpha, p_2^*) - \sum_{x\in S_1} b(p_2^*)(x)\sum_{y\in S_2} p_2^*(y)U_1(\alpha,y) \\
 & = & U_1(\alpha, p_2^*) - U_1(\alpha, p_2^*)\cdot\sum_{x\in S_1} b(p_2^*)(x)
\end{array}
$$

$$
\begin{aligned}
\overline{U}_1(b(p_2^*), p_2^*) &= \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) \overline{u}_1(x, y) \\
&= \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) [u_1(x, y) - U_1(\alpha, y)] \\
&= \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) u_1(x, y) \\
&- \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) U_1(\alpha, y) \\
&= U_1(b(p_2^*), p_2^*) \\
&- \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) U_1(\alpha, y) \\
&> U_1(\alpha, p_2^*) - \textstyle\sum_{x \in S_1} \sum_{y \in S_2} b(p_2^*)(x) p_2^*(y) U_1(\alpha, y) \\
&= U_1(\alpha, p_2^*) - \textstyle\sum_{x \in S_1} b(p_2^*)(x) \sum_{y \in S_2} p_2^*(y) U_1(\alpha, y) \\
&= U_1(\alpha, p_2^*) - {\color{blue} U_1(\alpha, p_2^*) \cdot \textstyle\sum_{x \in S_1} b(p_2^*)(x)} \\
&= U_1(\alpha, p_2^*) - {\color{blue} U_1(\alpha, p_2^*)} = 0
\end{aligned}
$$

Hence, for all $m \in \Delta(S_2)$ we have

$$\overline{U}(p_1^*, m) \geq \overline{U}_1(p_1^*, p_2^*) \geq \overline{U}_1(b(p_2^*), p_2^*) > 0$$

Hence, for all $m \in \Delta(S_2)$ we have

$$\overline{U}(p_1^*, m) \geq \overline{U}_1(p_1^*, p_2^*) \geq \overline{U}_1(b(p_2^*), p_2^*) > 0$$

which implies for all $m \in \Delta(S_2)$, $U_1(p_1^*, m) > U_1(\alpha, m)$, and so α is strictly dominated by p_1^*.

Important Issue: Correlated Beliefs

x	l	r
u	1,1,3	1,0,3
d	0,1,0	0,0,0

y	l	r
u	1,1,2	1,0,0
d	0,1,0	1,1,2

z	l	r
u	1,1,0	1,0,0
d	0,1,3	0,0,3

Important Issue: Correlated Beliefs

x	l	r
u	1,1,3	1,0,3
d	0,1,0	0,0,0

y	l	r
u	1,1,2	1,0,0
d	0,1,0	1,1,2

z	l	r
u	1,1,0	1,0,0
d	0,1,3	0,0,3

- Note that y is not strictly dominated for Charles.

Important Issue: Correlated Beliefs

x	l	r
u	1,1,3	1,0,3
d	0,1,0	0,0,0

y	l	r
u	1,1,2	1,0,0
d	0,1,0	1,1,2

z	l	r
u	1,1,0	1,0,0
d	0,1,3	0,0,3

- Note that y is not strictly dominated for Charles.
- It is easy to find a probability measure $p \in \Delta(S_A \times S_B)$ such that y is a best response to p. Suppose that $p(u, l) = p(d, r) = \frac{1}{2}$. Then, $EU(x, p) = EU(z, p) = 1.5$ while $EU(y, p) = 2$.

Important Issue: Correlated Beliefs

x	l	r
u	1,1,3	1,0,3
d	0,1,0	0,0,0

y	l	r
u	1,1,2	1,0,0
d	0,1,0	1,1,2

z	l	r
u	1,1,0	1,0,0
d	0,1,3	0,0,3

- Note that y is not strictly dominated for Charles.
- It is easy to find a probability measure $p \in \Delta(S_A \times S_B)$ such that y is a best response to p. Suppose that $p(u, l) = p(d, r) = \frac{1}{2}$. Then, $EU(x, p) = EU(z, p) = 1.5$ while $EU(y, p) = 2$.
- However, there is no probability measure $p \in \Delta(S_A \times S_B)$ such that y is a best response to p and $p(u, l) = p(u) \cdot p(l)$.

x	l	r
u	1,1,3	1,0,3
d	0,1,0	0,0,0

y	l	r
u	1,1,2	1,0,0
d	0,1,0	1,1,2

z	l	r
u	1,1,0	1,0,0
d	0,1,3	0,0,3

- To see this, suppose that a is the probability assigned to u and b is the probability assigned to l. Then, we have:
  - The expected utility of y is $2ab + 2(1-a)(1-b)$;
  - The expected utility of x is $3ab + 3a(1-b) = 3a(b + (1-b)) = 3a$; and
  - The expected utility of z is $3(1-a)b + 3(1-a)(1-b) = 3(1-a)(b + (1-b)) = 3(1-a)$.

Weak Dominance and MEU

Fact. Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N}\rangle$ is a strategic game and $X \subseteq S_{-i}$. A strategy $s_i \in S_i$ is weakly dominated (possibly by a mixed strategy) with respect to X iff there is **no full support probability measure** $p \in \Delta^{>0}(X)$ such that s_i is a best response to p.

Model of Differential Information

- Let Ω be a set of states:
  "The term "state of the world" implies a definite specification of all parameters that may be the object of uncertainty on the part of any player of G. In particular, each w includes a specification of which action is chosen by each player of G at that state w. Conditional on a given world, everybody knows everything; but in general, nobody knows which is really the true w." (pg. 6)

Model of Differential Information

- Let Ω be a set of states:
  "The term "state of the world" implies a definite specification of all parameters that may be the object of uncertainty on the part of any player of G. In particular, each w includes a specification of which action is chosen by each player of G at that state w. Conditional on a given world, everybody knows everything; but in general, nobody knows which is really the true w." (pg. 6)

- Let $\{\Pi_i\}_{i \in N}$ be a set of partitions on Ω, $\Pi_i(w)$ is the element of Π_i containing w

Model of Differential Information

- Let Ω be a set of states:
  "The term "state of the world" implies a definite specification of all parameters that may be the object of uncertainty on the part of any player of G. In particular, each w includes a specification of which action is chosen by each player of G at that state w. Conditional on a given world, everybody knows everything; but in general, nobody knows which is really the true w." (pg. 6)

- Let $\{\Pi_i\}_{i\in N}$ be a set of partitions on Ω, $\Pi_i(w)$ is the element of Π_i containing w

- Let $\{p_i\}_{i\in N}$ is a set of probability measure on Ω

Common Prior Assumption

Common Prior Assumption (CPA): There is a probability measure p on Ω such that

$$p_1 = p_2 = \cdots = p_n = p$$

Prior and posterior beliefs

- Different stages of information disclosure:

Prior and posterior beliefs

- Different stages of information disclosure:

  1. *Ex ante*: No information. Possibly background beliefs about who you are playing against (background beliefs about the "context of interaction").

Prior and posterior beliefs

- Different stages of information disclosure:

  1. *Ex ante*: No information. Possibly background beliefs about who you are playing against (background beliefs about the "context of interaction").

  2. *Ex interim*: The players receive a (private) information (a "signal". They know more.

Prior and posterior beliefs

- Different stages of information disclosure:

  1. *Ex ante*: No information. Possibly background beliefs about who you are playing against (background beliefs about the "context of interaction").
     Prior beliefs.

  2. *Ex interim*: The players receive a (private) information (a "signal". They know more.
     Posterior beliefs.

Common priors

- $\not\Rightarrow$ same posteriors!

Common priors

- $\not\Rightarrow$ same posteriors! For the simple reason that agents can receive different private information.

Common priors

- $\not\Rightarrow$ same posteriors! For the simple reason that agents can receive different private information.
- We play card together.

Common priors

- $\not\Rightarrow$ same posteriors! For the simple reason that agents can receive different private information.
- We play card together. Before the cards are dealt, our common prior belief that the other end up with a Joker is $0.037 = 2/54$.

Common priors

- $\not\Rightarrow$ same posteriors! For the simple reason that agents can receive different private information.
- We play card together. Before the cards are dealt, our common prior belief that the other end up with a Joker is $0.037 = 2/54$.
- We get 5 card each (and don't show them to each other). I end up with the 2 Jokers.
  - My posterior belief that you have a Joker is 0.
  - Your posterior belief that I have a Joker is $0.04 = 2/49$.

Common Prior Assumption

- Differences in posterior beliefs should be seen as coming from different information, not from different priors.

Common Prior Assumption

- Differences in posterior beliefs should be seen as coming from different information, not from different priors.
- Alternative formulation: "...people with different information may legitimately entertain different probabilities, but there is no rational basis for people who have always been fed precisely the same information to do so." (...Given that everyone are Bayesians)

Common Prior Assumption

- Differences in posterior beliefs should be seen as coming from different information, not from different priors.
- Alternative formulation: "...people with different information may legitimately entertain different probabilities, but there is no rational basis for people who have always been fed precisely the same information to do so." (...Given that everyone are Bayesians)

"Harsanyi doctrine" [Aumann, 1976].

R. Aumann. *Agreeing to Disagree*. *Annals of Statistics*, Vol.4, No.6, 1976.

J.C. Harsanyi. *Games with incomplete informations played by bayesian players*. *Management Science* 14:159182, 320334, 486502, 1967-68.

Common Prior Assumption

- Differences in posterior beliefs should be seen as coming from different information, not from different priors.
- Alternative formulation: "...people with different information may legitimately entertain different probabilities, but there is no rational basis for people who have always been fed precisely the same information to do so." (...Given that everyone are Bayesians)

"Harsanyi doctrine" [Aumann, 1976].

- CPA not an innocuous assumption! (cf. Aumann's agreeing to disagree theorem)

R. Aumann. *Agreeing to Disagree*. *Annals of Statistics*, Vol.4, No.6, 1976.

J.C. Harsanyi. *Games with incomplete informations played by bayesian players*. *Management Science* 14:159182, 320334, 486502, 1967-68.

Why CPA?

Why CPA?

- A controversial question!
  - See Morris (1995) for a thorough discussion.

Why CPA?

- A controversial question!
  - See Morris (1995) for a thorough discussion. One methodological observation:

Why CPA?

- A controversial question!
  - See Morris (1995) for a thorough discussion. One methodological observation:
    - Better to explain differences in posterior on the basis of identifiable differences in information or plausible errors in information processing.

Why CPA?

- A controversial question!
  - See Morris (1995) for a thorough discussion. One methodological observation:
    - Better to explain differences in posterior on the basis of identifiable differences in information or plausible errors in information processing.
    - Resorting on differences in priors often appears ad hoc (the resulting theory is "too permissive").

S. Morris. *The Common Prior Assumption in Economic Theory*. *Economics and Philosophy*, 11(2): pgs. 227- 253, 1995.

Let $G = \langle \{S_i\}_{i\in N}, \{u_i\}_{i\in N}\rangle$ be a strategic game.

- Ω is a set of states
- There is a common prior: a probability measure p on Ω
- $\{\Pi_i\}_{i\in N}$ is the set of information partitions
- $\mathbf{s} : \Omega \rightarrow S_1 \times \cdots \times S_n$ with $\mathbf{s}_i(w)$ the strategy of player i at w.
- For each $i \in N$, the players "know" which action she chooses: $\mathbf{s}_i$ is measurable with respect to Π_i:
- The expected utility of the strategy choice at w is:

$$EU_i(\mathbf{s}, w) = \sum_{v\in\Omega} p(v \mid \Pi_i(w))u_i(\mathbf{s}(v))$$

Let $G = \langle \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a strategic game.

- Ω is a set of states
- There is a common prior: a probability measure p on Ω
- $\{\Pi_i\}_{i \in N}$ is the set of information partitions
- $\mathbf{s} : \Omega \to S_1 \times \cdots \times S_n$ with $\mathbf{s}_i(w)$ the strategy of player i at w.
- For each $i \in N$, the players "know" which action she chooses: $\mathbf{s}_i$ is measurable with respect to Π_i:
- The expected utility of the strategy choice at w is:

$$EU_i(\mathbf{s}, w) = \sum_{v \in \Omega} p(v \mid \Pi_i(w)) u_i(\mathbf{s}(v))$$

Bayes Rationality For all $w \in \Omega$, $EU_i(\mathbf{s}, w) \geq EU((s_i, \mathbf{s}_{-i}), w)$ for all $s_i \in S_i$

Correlated Equilibrium

Let $G = \langle \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a game.

A correlated strategy n-tuple in G is a function from a finite probability space Γ into $S = S_1 \times \cdots \times S_n$. That is, f is a *random variable* whose values are n-tupels of actions.

Chance (according to the probability space Γ) chooses an element $\gamma \in \Gamma$, then each player is recommended to take action $f_i(\gamma)$.

Correlated Equilibrium: A correlated equilibrium in G is a correlated strategy n-tuple f such that

$$Eu_i(f) \geq Eu_i(g_i, f_{-i})$$

Theorem. Assume that there is a common prior and that for all w, for all $i \in N$, $\Pi_i(w) \subseteq \{v \mid \mathbf{s}_i(v) = \mathbf{s}_i(w)\}$. If each player is Bayes rational at each state of the world, then the distribution of the action n-tuple $\mathbf{s}$ is a correlated equilibrium.

Nash Equilibrium

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

- The profiles **aA** and **bB** are two pure-strategy Nash equilibria of that game.

Definition

A strategy profile σ is a *Nash equilibrium* iff for all i and all $s_i' \neq \sigma_i$:

$$u_i(\sigma) \geq u_i(s_i, \sigma_{-i})$$

More Specific Expectations

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

More Specific Expectations

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

- If Ann believes that Bob plays **A**, the only rational choice for her is **a**.

More Specific Expectations

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

- If Ann believes that Bob plays **A**, the only rational choice for her is **a**.
- The same hold for Bob.

More Specific Expectations

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

- If Ann believes that Bob plays **A**, the only rational choice for her is **a**.
- The same hold for Bob.
- If, furthermore, these beliefs are *true*, then **aA** is played.

Knowledge of Strategies and Nash Equilibrium

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

- If Ann and Bob are rational and have correct beliefs about each others' strategy choices, then **aA** is played.

Knowledge of Strategies and Nash Equilibrium

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

- If Ann and Bob are rational and have correct beliefs about each others' strategy choices, then **aA** is played.
- For any two-players strategic game and model for that game, if at state w both players are rational and know the other's strategy choice, then $\sigma(w)$ is a Nash equilibrium.

R. Aumann and A. Brandenburger, "Epistemic Conditions for Nash Equilibrium". *Econometrica*. 1995.

Hard Knowledge of Strategies and Nash Equilibrium

Theorem

(Aumann and Brandenburger, 1995) For any two-players strategic game and model for that game, if at state w both players are rational and know other's strategy choice, then $\sigma(w)$ is a Nash equilibrium.

- Remarks:

Hard Knowledge of Strategies and Nash Equilibrium

Theorem

(Aumann and Brandenburger, 1995) For any two-players strategic game and model for that game, if at state w both players are rational and know other's strategy choice, then $\sigma(w)$ is a Nash equilibrium.

- Remarks:
  - Close to the intuitive explanation: Best response *given* the choices of others, or no regret.

Hard Knowledge of Strategies and Nash Equilibrium

Theorem

(Aumann and Brandenburger, 1995) For any two-players strategic game and model for that game, if at state w both players are rational and know other's strategy choice, then $\sigma(w)$ is a Nash equilibrium.

- Remarks:
  - Close to the intuitive explanation: Best response *given* the choices of others, or no regret.
  - No higher-order information needed... for 2 players (more on this in a moment).

Hard Knowledge of Strategies and Nash Equilibrium

Theorem

(Aumann and Brandenburger, 1995) For any two-players strategic game and model for that game, if at state w both players are rational and know other's strategy choice, then $\sigma(w)$ is a Nash equilibrium.

- Remarks:
  - Close to the intuitive explanation: Best response *given* the choices of others, or no regret.
  - No higher-order information needed... for 2 players (more on this in a moment).
  - Hard knowledge, or even correct beliefs, about actions taken? Does Nash equilibrium undermine strategic uncertainty?

Nash equilibrium, the general case

(Aumann and Brandenburger, 1995) In an n-player game, suppose that the players have a common prior, that their payoff functions and their rationality are mutually known, and that their conjectures are commonly known. Then for each player j, all the other players i agree on the same conjecture σ_j about j, and the resulting profile $(\sigma_1, .., \sigma_n)$ of mixed actions is a Nash equilibrium.

- Remarks:

Nash equilibrium, the general case

(Aumann and Brandenburger, 1995) In an n-player game, suppose that the players have a common prior, that their payoff functions and their rationality are mutually known, and that their conjectures are commonly known. Then for each player j, all the other players i agree on the same conjecture σ_j about j, and the resulting profile $(\sigma_1, .., \sigma_n)$ of mixed actions is a Nash equilibrium.

- Remarks:
  - Higher-order information after all: common knowledge of conjectures.

Nash equilibrium, the general case

(Aumann and Brandenburger, 1995) In an n-player game, suppose that the players have a common prior, that their payoff functions and their rationality are mutually known, and that their conjectures are commonly known. Then for each player j, all the other players i agree on the same conjecture σ_j about j, and the resulting profile $(\sigma_1, .., \sigma_n)$ of mixed actions is a Nash equilibrium.

- Remarks:
  - Higher-order information after all: common knowledge of conjectures.
  - The result is "tight". Fails if we drop any of the conditions.

Nash equilibrium, the general case

(Aumann and Brandenburger, 1995) In an n-player game, suppose that the players have a common prior, that their payoff functions and their rationality are mutually known, and that their conjectures are commonly known. Then for each player j, all the other players i agree on the same conjecture σ_j about j, and the resulting profile $(\sigma_1, .., \sigma_n)$ of mixed actions is a Nash equilibrium.

- Remarks:
  - Higher-order information after all: common knowledge of conjectures.
  - The result is "tight". Fails if we drop any of the conditions.
  - Epistemic Interpretation of mixed strategies.

Nash equilibrium, the general case

(Aumann and Brandenburger, 1995) In an n-player game, suppose that the players have a common prior, that their payoff functions and their rationality are mutually known, and that their conjectures are commonly known. Then for each player j, all the other players i agree on the same conjecture σ_j about j, and the resulting profile $(\sigma_1, .., \sigma_n)$ of mixed actions is a Nash equilibrium.

- Remarks:
  - Higher-order information after all: common knowledge of conjectures.
  - The result is "tight". Fails if we drop any of the conditions.
  - Epistemic Interpretation of mixed strategies.
  - If the payoffs are common knowledge, then rationality is also common knowledge (Ben Polak, Econometrica, 1999).

Nash equilibrium, the general case

(Aumann and Brandenburger, 1995) In an n-player game, suppose that the players have a common prior, that their payoff functions and their rationality are mutually known, and that their conjectures are commonly known. Then for each player j, all the other players i agree on the same conjecture σ_j about j, and the resulting profile $(\sigma_1, .., \sigma_n)$ of mixed actions is a Nash equilibrium.

- Remarks:
  - Higher-order information after all: common knowledge of conjectures.
  - The result is "tight". Fails if we drop any of the conditions.
  - Epistemic Interpretation of mixed strategies.
  - If the payoffs are common knowledge, then rationality is also common knowledge (Ben Polak, Econometrica, 1999).
  - But still, CKR does not imply Nash Equilibrium.

The Importance of Correlations

The Importance of Correlations

- We will see that rationality and common belief of rationality implies that players play correlated rationalizble strategies.

The Importance of Correlations

- We will see that rationality and common belief of rationality implies that players play correlated rationalizble strategies.
- Question: where do (beliefs in) correlations come from?
  Answer: A player can think that other players' strategy choices are correlated, because he thinks what they believe about the game is correlated.

Two routes to explain correlations

Correlations

Two routes to explain correlations

Extrinsic

Correlations

R.J. Aumann. *Correlated equilibrium as an expression of bayesian rationality*. *Econometrica*, 55(1-18), 1987.

Two routes to explain correlations

Signals

Correlations

R.J. Aumann. *Correlated equilibrium as an expression of bayesian rationality*. *Econometrica*, 55(1-18), 1987.

Two routes to explain correlations

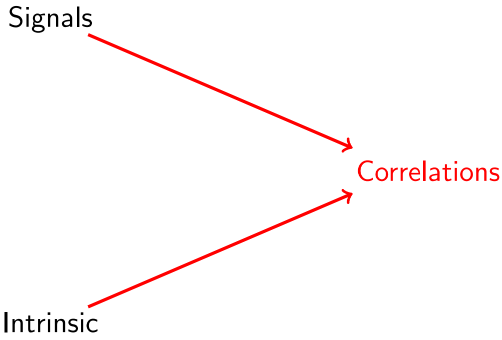

A. Brandenburger and A. Friedenberg. *Intrinsic Correlations in Games*. J.E.T., vol. 141, 2008.

Two routes to explain correlations

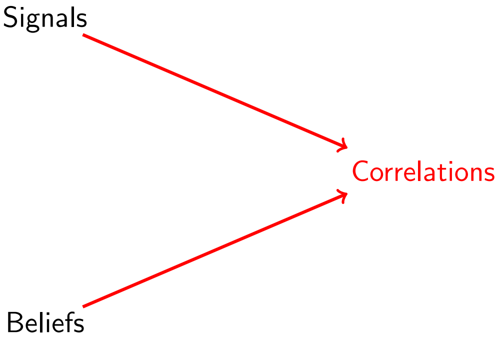

A. Brandenburger and A. Friedenberg. *Intrinsic Correlations in Games*. J.E.T., vol. 141, 2008.

Two routes to explain correlations

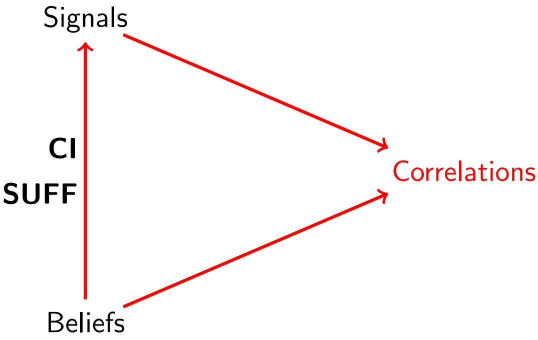

R.J. Aumann. *Correlated equilibrium as an expression of bayesian rationality*. *Econometrica*, 55(1-18), 1987.

A. Brandenburger and A. Friedenberg. *Intrinsic Correlations in Games*. J.E.T., vol. 141, 2008.

Two routes to explain correlations

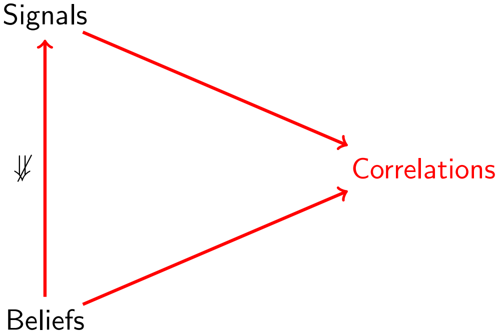

A. Brandenburger and A. Friedenberg. *Intrinsic Correlations in Games*. J.E.T., vol. 141, 2008.