1 Normal Form Games

A normal form game is $(I, (A^i)_{i=1,\dots,n}, (u^i)_{i=1,\dots,n})$, where $\forall i$ A_i is an action set, $A = \times_{i=1}^n A^i$, and $u_i : A \to \mathbb{R}$ $\forall i$.

$I = \{1, ..., n\}$ is the set of players.

Assume A^i is finite for all i.

Examples: Coordination, Matching pennies, Prisoner's dilemma, Battle of the Sexes.

1.1 Dominance

Let $S^i = \Delta(A^i) = \{(s(a_1^i), ..., s(a_{k_i}^i)) : \forall i,\ s(a_i) \geq 0, \sum_{A^i} s(a^i) = 1\}$.

A mixed extension of a normal form game is $(I, (S^i)_{i=1,\dots,n}, (u^i)_{i=1,\dots,n})$, where $\forall$ $S^i = \Delta(A^i)$, $S = \prod_{i=1}^n S^i$ and $u^i : S \to \mathbb{R}$ is defined by

$$u^i(s^1, ..., s^n) = \sum_{a \in A} u^i(a) \prod_{i=1}^n s^i(a^i).$$

We write $Pr_s(a) = \prod_{i=1}^n s^i(a^i) \in \Delta A$.

Example: Show that in the game below, the player can get a better payoff by mixing T and M than by playing B, no matter what his belief is about what his partner is doing.

	L	R
T	3	0
M	0	3
B	1	1

We say $s^i \in S^i$ strictly dominates $a^i \in A^i$ iff for all a^{-i}

$$u^i(s^i, a^{-i}) > u^i(a^i, a^{-i}).$$

Alternatively,

$$s^i D_2 a^i \Leftrightarrow \forall s^{-i} \in S^{-i} \quad u^i(s^i, s^{-i}) > u_i(a^i, s^{-i})$$

or

$$s^i D_3 a^i \Leftrightarrow \forall \mu \in \Delta(A^{-i}) \quad u^i(s^i, \mu) > u_i(a^i, \mu)$$

Exercise: $s^i D_3 a^i \Leftrightarrow s^i D_2 a^i \Leftrightarrow s^i D_1 a^i$.

Example: Note that T and L are both dominated in the game below.

	L	R
T	-2,-2	-10,-1
B	-1,-10	-5,-5

This leads to the counter-intuitive prediction of playing (B,R). Of course this doesn't happen in real life.

Example:

	L	R
T	3	0
M	0	3
B	x	x

Consider a belief p for Player 1 that Player 2 chooses L. Note that if $x < \frac{3}{2}$, B is never a best response. For every belief, Player 1 is better of playing T or M. Dually, $\exists s^1 \in S^1$ that dominates B.

If $x = \frac{3}{2}$, there exists a belief ($p = 0.5$) for which B is a best response. Dually, B is not strictly dominated.

This example suggests that an action is never a best response if and only if it is strictly dominated by a strategy.

Definition: An action $a^i \in A^i$ is never a best response if there is no $\mu \in \Delta(A^{-i})$ such that $u^i(a^i, \mu) \geq u^i(b^i, \mu)$ for all b^i.

Theorem: An action $a^i \in A^i$ is strictly dominated if and only if it is never a best response.

One direction is easy to prove (see your class notes). The proof for the other direction can be found in Osborne and Rubinstein.

1.2 IESDA

We illustrate this with examples:

	L	R
T	0,-2	-10,-1
B	-1,-10	-5,-5

	L	R
T	3,0	0,1
M	0,0	3,1
B	1,1	1,0

Formally, an iterated elimination of strictly dominated actions is a sequence $A_1, A_2, A_3, ..., A_T$ where:

1. $\forall k = 1, ..., T \quad A_k = \times_{i=1}^{N} A_k^i$.
2. $\forall k, i \quad A_{k+1}^i \subset A_k^i$
3. $\forall i \quad A_1^i = A^i$
4. Elimination rule: $\forall k, i \quad a^i \in A_k^i \setminus A_{k+1}^i$ only if $\exists s^i \in \Delta(A_k^i)$ s.t. $\forall a^{-i} \in A_k^{-i} \quad u^i(s^i, a^{-i}) > u^i(a^i, a^{-i})$.

Example (Cournot Duopoly):

Consider a two player game with two firms $i = 1, 2$. Each firm faces the demand curve $p = a - b(q_1 + q_2)$ and per-unit costs of production c. Show that iterated elimination of strictly dominated actions yields a unique outcome in which each firm produces $\frac{a-c}{3b}$.

1.3 Rationalizability

An action $a^i \in A^i$ is rationalizable if there exists a vector $(R^1, ..., R^N)$ such that:

1. $a^i \in R^i$
2. For all j, $R^j \subset A^j$

3. $\forall j, b_j \in R^j, \quad \exists \mu(b^j) \in \Delta(A^{-j})$ (with support R^{-j}) s.t. $u^j(b^j, \mu) \geq u^j(a^j, \mu) \quad \forall a^j \in A^j$.

We will also sometimes talk of sets as being rationalizable. In this case, we will talk of $(R^1, ..., R^N)$ as being rationalizable if conditions 2 and 3 above are satisfied.

Example:

	L	R
T	3,0	0,1
M	0,0	3,1
B	1,1	1,0

$(R^1, R^2) = (\{M\}, \{R\})$.

Example 2:

	L	R
T	3,1	0,0
B	0,0	1,3

$(R^1, R^2) = (\{T\}, \{L\})$. $(T^1, T^2) = (\{B\}, \{R\})$.

So these sets are not unique, unless maximality is required. The maximal set of rationalizable actions in this example is $\{\{T, B\}, \{L, R\}\}$.

Proposition: If $(R^1, ..., R^N)$ and $(T^1, ..., T^N)$ are rationalizable, then $(R^1 \cup T^1, ..., R^N \cup T^N)$ is rationalizable, as well.

As we discussed before, $D \Leftrightarrow NBR$. D is related to iterated dominance, while NBR is related to rationalizability. As we will show below, IESDA and rationalizability are in some sense equivalent:

$$IESDA \Leftrightarrow Rationalizability$$

We will see this from two theorems, the first of which is below:

Theorem: Let R be a set of rationalizable actions. Let $(A_1, ..., A_T)$ be an iterated elimination of strictly dominated actions. Then, $R^i \subset A^i_T \quad \forall i$.

Proof: The proof proceeds by induction of T, the number of steps in the elimination procedure.

1 (initial step): $R^i \subset A_1^i = A^i$ by definition.

2. (inductive step): Assume $R^i \subset A_n^i$. Let $a^i \in R^i$. Then,

$$\exists \mu_{a^i} \in \Delta(A^{-i}) \text{ (with support } R^{-i}) \quad s.t. \quad u^i(a^i, \mu_{a^i}) \geq u^i(b^i, \mu_{a^i}) \quad \forall b^i \in A^i.$$

Since $R^i \subset A_n^i$,

$$\exists \mu_{a^i} \in \Delta(A_n^{-i}) \quad s.t. \quad u^i(a^i, \mu_{a^i}) \geq u^i(b^i, \mu_{a^i}) \quad \forall b^i \in A^i.$$

By a theorem we already now, $D \Rightarrow NBR$. An equivalent of stating this is $BR \rightarrow$ not D. Hence,

$$\nexists s^i \in \Delta(A^i) \quad s.t. \quad u^i(s^i, a^{-i}) > u^i(a^i, a^{-i}) \quad \forall a^{-i} \in A_n^{-i}.$$

I.e., a^i is not eliminated in the n^{th} stage of the elimination procedure. Hence, $a^i \in A_{n+1}^i$, as needed.

Theorem 2: Let R denote the maximal set of rationalizable actions (in terms of set inclusion), and let $(A_1, ..., A_T)$ be a complete elimination of strictly dominated strategies. Then $A_T^i \subset R^i$ for every i.

Proof: It enough to show that $(A_T^1, ..., A_T^N)$ is rationalizable.

Let $a^i \in A_T^i$. Then $\nexists s^i \in \Delta(A_T^i) \quad s.t. \quad u^i(s^i, a^{-i}) > u^i(a^i, a^{-i})$ for all $a^{-i} \in A_T^{-i}$.

Note that this implies the same statement with $s^i \in \Delta(A_T^i)$ replaced by $s^i \in \Delta(A^i)$.

To see this, let $s^i \in \Delta(A^i)$ be some strategy that puts a positive probability on an action d^i that was removed at an earlier step. Since the action was removed, you can find some $e^i \in A_T^i$ s.t. $u^i(e^i, a^{-i}) > u^i(d^i, a^{-i})$ for all $a^{-i} \in A_T^{-i}$. Transfer the probability placed on d^i onto e^i. Do the same for all other actions that didn't survive the elimination procedure to obtain a mixed strategy $\tilde{s}^i$ with support A_T^i. Note that $u^i(\tilde{s}^i, a^{-i}) > u^i(s^i, a^{-i})$ for all $a^{-i} \in A_T^{-i}$. Since you can't find a $\tilde{s}^i$ that dominates a^i in the last step of the elimination procedure, you will not be able to find a s^i, either.

Next, recall that $\neg D \Rightarrow BR$. Therefore, $\exists \mu_{a^i} \in \Delta(A_T^{-i})$ such that $u^i(a^i, \mu_{a^i}) \geq u^i(b^i, \mu_{a^i}) \quad \forall b^i \in A^i$.

This shows that $(A^1_T, ..., A^N_T)$ is rationalizable, which is what we needed.

What the previous two theorems show is that that the set of actions that survives *complete* IESDA is unique and equal to the maximal set of rationalizable actions.

What the previous two theorems show is that that the set of actions that survives *complete* IESDA is unique and equal to the maximal set of rationalizable actions.

Example:

	b_1	b_2	b_3	b_4
a_1	0,7	2,5	7,0	0,1
a_2	5,2	3,3	5,2	0,1
a_3	7,0	2,5	0,7	0,1
a_4	0,0	0,-2	0,0	10,-1

It's easy to show that the maximal set of rationalizable actions is $(R^1, R^2) = (\{a_1, a_2, a_3\}, \{b_1, b_2, b_3\})$ (eliminate b_4 in step 1, and a_4 in step 2).

1.4 IEWDA

Definition: $s^i W a^i \quad \Leftrightarrow$

- For all $a^{-i} \in A^{-i} \quad u^i(s_i, a^{-i}) \geq u^i(a_i, a^{-i})$.
- There exists $b^{-i} \in A^{-i}$ such that $u^i(s_i, b^{-i}) > u^i(a_i, b^{-i})$.

Example: Using the game below, show that the set of actions surviving IEWDA is not unique.

	L	R
T	1,1	0,0
M	1,1	2,1
B	0,0	2,1

Example (Vickrey auction): An object is auctioned off to N bidders. Bidder i's valuation of the object (in monetary terms) is v_i. The auction rules are that each player submit a bid (a nonnegative number) in a sealed envelope. The envelopes are then opened, and the bidder who has submitted the highest bid gets the object but pays the auctioneer the amount of the *second-highest* bid. If more than one bidder submits the highest bid, each gets the object with equal probability. Show that submitting a bid of v_i with certainty is a weakly dominant strategy for bidder i.

1.5 Nash Equilibrium

Definition (Pure Best Reply): $PBR^i(s) = \{a^i \in A^i : u^i(a^i, s^{-i}) \geq u^i(b^i, s^{-i}) \quad \forall b^i \in A^i\}$.

Note that this set is nonempty for a finite game.

Definition (Best Reply): $BR^i(s) = \{s^i \in S^i : u^i(s^i, s^{-i}) \geq u^i(b^i, s^{-i}) \quad \forall b^i \in A^i\}$.

Note that for every s, $BR^i(s)$ is closed, convex, nonempty, and equal to the mixed strategies concentrated on $PBR^i(s)$.

Example: Find $PBR^i(s)$ and $BR^i(s)$ in the following game:

	L	R
T	3,1	0,0
B	0,0	1,3

Definition:

$$BR(s) = \times_{i=1}^{N} BR^i(s).$$

Note that $BR : S \twoheadrightarrow S$ is a closed, convex, and nomempty valued correspondence.

Mathematical Preliminaries for Existence of Nash

The proof of existence of Nash Equilibrium for finite games (what got Nash the Nobel prize) is a simple application of Kakutani's Fixed Point Theorem. These notes will provide the bare minimum necessary for understanding this theorem. For details, see Mas-Colell, Whinston, and Greene.

Definition (Correspondence): Given sets $A \subset \mathbb{R}^N$ and $Y \subset \mathbb{R}^K$, a correspondence $f : A \twoheadrightarrow Y$ is a rule that assigns a set $f(x) \subset Y$ to every $x \in A$.

Definition (Upperhemicontinuity): Given a set $A \subset \mathbb{R}^N$ and a **compact** $Y \subset$

$\mathbb{R}^K$, a correspondence $f : A \twoheadrightarrow Y$ is upperhemicontinuous if for any two sequences x_n and y_n s.t. $x_n \in A \ \forall n$, $y_n \in f(x_n)$ for all n, $x_n \to x \in A$, $y_n \to y \in Y$, it is also the case that $y \in f(x)$.

Roughly speaking, you can think of upperhemicontunuity as a generalization of the continuity notion for functions to correspondences.

Theorem (Kakutani's Fixed Point Theorem): Suppose that $A \subset \mathbb{R}^N$ is a nonempty, compact, convex set and that $f : A \twoheadrightarrow A$ an upperhemicontinuous correspondence with the property that for each $x \in A$, $f(x) \subset A$ is nonempty and convex. Then $f(\cdot)$ has a fixed point. I.e., there exists an $x \in A$ s.t. $x \in f(x)$.

Definition (Nash Equilibrium): A Nash Equilibrium of a NFG is a strategy profile $\hat{s}$ s.t. $\hat{s} \in BR(\hat{s})$.

Theorem: The set of Nash Equilibrium strategy profiles is nonempty.

Proof: The set S is nonempty, compact, and convex. Also, $BR : S \twoheadrightarrow S$ is convex and nonempty valued. It remains to show that $BR(\cdot)$ is upperhemicontinuous. For this, it is sufficient to show that each $BR^i(\cdot)$ is upperhemicontinuous.

Take a sequence s_n such that $s_n \in S \ \forall n$ and a sequence t_n such that

- $t_n^i \in BR^i(s_n) \ \forall i, n$;
- $s_n \to s \in S$;
- $t_n^i \to t^i$ for all i.

We need to show that $t^i \in BR^i(s)$ for all i. This is a simple consequence of contunuity of u. In particular, we know that

$$u^i(t_n^i, s_n^{-i}) \geq u^i(a^i, s_n^{-i}) \quad \forall a^i \in A^i, \ \forall n.$$

Take $\lim_{n\to\infty}$ of both sides to get:

$$u^i(t^i, s^{-i}) \geq u^i(a^i, s^{-i}) \quad \forall a^i \in A^i.$$

Philosophical point: Keep in mind that Nash equilibrium is *not* an implication of rationality. Rationalizability is an implication of rationality. Nash has stronger "epistemic" assumptions. E.g., it assumes that every player knows what every other player is playing.

1.6 Perfect Equilibrium

The idea of perfect equilibrium (PE) is motivated by the possibility that nobody's perfect (people make mistakes). As motivation, consider the following game:

	L	R
T	1,1	0,0
B	0,0	0,0

This game has two NE in pure strategies, (T,L) and (B,R), and no other mixed or pure strategy equilibria. However, the equilibrium (B,R) is not robust to small mistakes. To see this, let ϵ be the smallest probability of any action being chosen. Player 1's payoff from any mixed strategy profile $((p, 1-p), (q, 1-q))$ is pq. When the possibility of mistakes is introduced, $q \geq \epsilon$. Thus, Player 1 chooses the maximum p he can, which is $1-\epsilon$. Likewise, Player 2 will put probability $(1-\epsilon)$ on L against any mixed strategy of Player 1.Thus, the only equilibrium of the perturbed game is $((1-\epsilon, \epsilon), (1-\epsilon, \epsilon))$. This equilibrium goes to (T,L)=$((1,0),(1,0))$ as the probability of making a mistake goes to 0. (T,L) is called a **perfect equilibrium**.

In general, given any $\epsilon = ((\epsilon_1^1, ..., \epsilon_{\#A^1}^1), ..., (\epsilon_1^N, ..., \epsilon_{\#A^N}^N))$, define

$$S_\epsilon^i = \{s^i \in S^i : s^i(a_j^i) \geq \epsilon_j^i \quad \forall j \in \{1, 2, ..., \#A^i\}\}.$$

Definition: Given any mixed extension of a normal form game G, the **perturbed game** G^ϵ consists of the same players and the same utility functions, but with S^i replaced by S_ϵ^i for every i.

Definition: The set of Nash Equilibria of the perturbed game is written as $NE(G^\epsilon)$ and defined as all $s \in S_\epsilon$ such that

$$u^i(s^i, s^{-i}) \geq u^i(t^i, s^{-i}) \quad \forall i, \forall t^i \in S_\epsilon^i.$$

Definition: A mixed strategy $s \in S$ is called a **perfect equilibrium** if there exist sequences $\{\epsilon(n)\}$ and $\{s(n)\}$ such that the following conditions are satisfied:

1. $\epsilon_j^i(n) > 0 \quad \forall i \in \{1, ..., N\}, j \in \{1, ..., \#A^i\}, n \in \{1, 2, 3, 4, ...\}$.

2. $\sum_j \epsilon^i_j(n) < 1 \quad \forall i, n.$

3. $\lim_{n\to\infty} \epsilon^i_j(n) = 0 \quad \forall i, j.$

4. $s(n) \in NE(G^{\epsilon(n)}) \quad \forall n.$

5. $\lim_{n\to\infty} s(n) = s.$

Theorem:

1. Given any $\epsilon = ((\epsilon^1_1, ..., \epsilon^1_{\#A^1}), ..., (\epsilon^N_1, ..., \epsilon^N_{\#A^N}))$, the set $NE(G^\epsilon)$ is nonempty.
2. The set of perfect equilibria is nonempty.
3. Any perfect equilibrium is a Nash equilibrium.

Proof:

1. Homework.
2. Let $\epsilon(n)$ be any sequence s.t. conditions 1-3 in the definition of perfect equilibrium are satisfied, e.g. $\epsilon(n) = \frac{1}{n} \quad \forall n$ works for large enough n. By part 1 of the theorem, $\exists s(n) \in NE(G^{\epsilon(n)}) \quad \forall n$. $s(n)$ is a sequence in S, a compact set, and hence has a convergent subsequence $s(m(n)) \to s \in S$. The limit of the sequence is a perfect equilibrium by definition.
3. Since s is a perfect equilibrium, we know that there exists $s(n)$ s.t.

$$u^i(s^i(n), s^{-i}(n)) \geq u^i(r^i, s^{-i}(n)) \quad \forall i,\ \forall r^i \in S^i_{\epsilon(n)}.$$

Now, fix i and let t^i be any action in S^i. Since $\epsilon(n) \to 0$, we can find a sequence $t^i(n) \to t^i$. Then,

$$u^i(s^i(n), s^{-i}(n)) \geq u^i(t^i(n), s^{-i}(n)).$$

Taking $lim_{n\to\infty}$ of both sides of the inequality we get

$$u^i(s^i, s^{-i}) \geq u^i(t^i, s^{-i}).$$

Since we can do this for every i and for any $t^i \in S^i$, s is a Nash Equilibrium.

1.7 Evolutionarily Stable Strategies

This is another refinement which applies evolutionary reasoning to game theoretic concepts. In particular, we now interpret players as animals or plants with strategies chosen by genetics, and players' utilities as *fitness*. Fitness can be thought of as the animal's survival potential.

The idea of being evolutionarily stable is the following. With some small probability ϵ, nature may introduce a mutant into a population of healthy agents playing some strategy b^*. We say that b^* is evolutionarily stable if it has greater survival potential than the mutant. In this case, the mutant will die out.

In particular, consider a population composed of $1-\epsilon$ players with strategy b^* and ϵ mutants with strategy b. b^* has greater fitness than b if

$$(1-\epsilon)u(b^*,b^*) + \epsilon u(b^*,b) > (1-\epsilon)u(b,b^*) + \epsilon u(b,b).$$

This condition is satisfied for some small ϵ if either of the following hold:

1. $u(b^*,b^*) > u(b,b^*)$;
2. $u(b^*,b^*) = u(b,b^*)$ and $u(b^*,b) > u(b,b)$.

Motivated by this, we introduce the following definition:

Definition (ESS): Let $G = (\{1,2\},\{B,B\},\{u^i\}_{i=1,2})$ be a symmetric game. $b^* \in B$ is an **evolutionarily stable strategy (ESS)** if $(b^*,b^*) \in NE(G)$ and $u(b^*,b) > u(b,b)$ for all $b \in BR(b^*)$.

Note: ESS is a *refinement* on NE, since it gives a more restrictive prediction.

Note 2: Any NE b^* with the property that no strategy other than b^* is a best response to b^* is an ESS. Nash equilibria with this property are called **strict** Nash equilibria.

Example: A non-strict NE may not be an ESS. Consider the $2x2$ symmetric game where $u(a,b) = 1 \quad \forall a,b$. Everything is a NE here, and nothing is ESS.

Example: Consider the following game, with $0 < \gamma < 1$:

	A	B	C
A	γ,γ	1,-1	-1, 1
B	-1, 1	γ,γ	1, -1
C	1,-1	-1, 1	γ,γ

(1/3,1/3,1/3) is the unique mixed strategy NE of this game. However, it is not evolutionary stable.

Example: Consider the following game:

0,0	2,2
2,2	0,0

The unique mixed NE of this game is ((1/2),(1/2)), and it is ESS.

Example (Hawk/Dove):

1/2,1/2	0,1
1,0	1/2(1-c), 1/2(1-c)

1.8 Correlated Equilibrium

The first step is to recognize that the set of probabilities induced on A by mixed strategies is not equal to $\Delta(A)$. Let $Pr_s(a) = \prod_{i=1}^{N} s^i(a^i)$. The claim is that

$$\{Pr_s(a)\} \subsetneq \Delta(A).$$

Example: Consider a $\mu(a)$ given by:

	L	R
T	$\frac{1}{2}$	0
B	0	$\frac{1}{2}$

The set $\{Pr_s(a)\}$ is:

	L	R
T	$s^1(T)s^2(L)$	$s^1(T)s^2(R)$
B	$s^1(B)s^2(L)$	$s^1(B)s^2(R)$

Clearly, $\mu(a) \notin \{Pr_s(a)\}$.

A **correlated equilibrium** is a probability distribution $\mu(a)$ such that if players are given recommendations based on the distribution, any player i finds it in his best interests to play a^i, assuming all other players do the same.

Example:

	L	R
T	3,1	0,0
B	0,0	1,3

Any $\mu(a) \in \Delta(A)$ is vector p s.t. $\sum p_{ij} = 1$, where ant $p_{ij} \geq 0$ is the probability assigned to the action a_{ij} in the payoff matrix:

	L	R
T	p_{11}	p_{12}
B	p_{21}	p_{22}

Assume that the players are given recommendations based on μ. Then,

- $\mu(\cdot|T) = (\frac{p_{11}}{p_{11}+p_{12}}, \frac{p_{12}}{p_{11}+p_{12}})$,
- $\mu(\cdot|B) = (\frac{p_{21}}{p_{21}+p_{22}}, \frac{p_{22}}{p_{21}+p_{22}})$,
- $\mu(\cdot|L) = (\frac{p_{11}}{p_{11}+p_{21}}, \frac{p_{21}}{p_{11}+p_{21}})$,
- $\mu(\cdot|R) = (\frac{p_{12}}{p_{12}+p_{22}}, \frac{p_{22}}{p_{12}+p_{22}})$.

It follows that

- Player 1 follows the recommendation T if $3 * p_{11} \geq 1 * p_{12}$,
- Player 1 follows the recommendation B if $1 * p_{22} \geq 3 * p_{21}$,
- Player 2 follows the recommendation L if $1 * p_{11} \geq 3 * p_{21}$,
- Player 2 follows the recommendation R if $3 * p_{22} \geq 1 * p_{12}$.

The set of correlated equilibria is the set of p's satisfying these inequalities. Note that any p derived from a Nash equilibrium works! Thus, the inequalities are satisfied by $(1, 0, 0, 0)$, $(0, 0, 0, 1)$, and $(\frac{3}{16}, \frac{9}{16}, \frac{1}{16}, \frac{3}{4})$. The question is if the set of correlated equilibria is bigger. And indeed it is. Clearly, any convex combination of the three correlated equilibria above is itself a correlated equilibrium.

Let's make this a bit more formal.

Definition (Correlated Equilibrium): $\mu \in \Delta(A)$ is a correlated equilibrium if $\forall i, a^i \in A^i, b^i \in A^i$,

$$\sum_{a^{-i}} u^i(a^i, a^{-i})\mu(a^{-i}|a^i) \geq \sum_{a^{-i}} u^i(b^i, a^{-i})\mu(a^{-i}|a^i).$$

Theorem: The set of correlated equilibria is non-empty, closed, and convex.

Proof: It is non-empty because any $Pr_s(a)$ derived from a Nash Equilibrium leads to a correlated equilibrium (check it). It is closed and convex because it is a system of linear inequalities. The details are left for homework.

Example (Game of Chicken):

	L	R
T	6,6	2,7
B	7,2	0,0

This game has three Nash equilibria, with payoffs (2,7), (7,2), and $(4\frac{2}{3}, 4\frac{2}{3})$. The μ below gives a correlated equilibrium with payoffs $(5, 5)$, which is outside the convex hull of Nash equilibrium payoffs.

	L	R
T	$\frac{1}{3}$	$\frac{1}{3}$
B	$\frac{1}{3}$	0

Example(Correlated Equilibrium in an Infinite Game)

$z < 0$:

	L	R
U	$2, 2, 3 + 1/z$	$0, 0, 8$
D	$0, 0, 0$	$2, 2, 0$

$z = 0$:

	L	R
U	$2, 2, 2$	$0, 0, 0$
D	$0, 0, 0$	$2, 2, 2$

$z > 0$:

	L	R
U	$2, 2, 0$	$0, 0, 0$
D	$0, 0, 8$	$2, 2, 3 - 1/z$

Here, $z \in \mathbb{Z}$, so player 3 has infinitely many actions.

This game has no pure or mixed strategy Nash Equilibria. To see this, notice that regardless of what player 3 is doing, players 1 and 2 will be playing the following game between themselves:

	L	R
U	$2, 2$	$0, 0$
D	$0, 0$	$2, 2$

The Nash Equilibria of this game are (U,L), (D,R), and a mixed strategy equilibrium where both players mix between their two actions with equal probabilities.

If (U,L) is played, 3 wants do decrease z to $-\infty$. If (D,R) is played, 3 wants to increase z to ∞. Similarly, in the mixed case Player 3 has no best response.

On the other hand, a correlated equilibrium exists. To see this, let $\mu(U, L, z = 0) = \mu(D, R, z = 0) = 1/2$.

1.9 Correlated Equilibrium as an Expression of Bayesian Rationality

Another important interpretation of correlated equilibria was provided by Aumann (1987). Assume we are exogenously given the following informational model of the world:

(i) A finite set Ω, with generic element ω,
(ii) A probability measure ρ on Ω,
(iii) $\forall i$, a partition $\mathcal{P}^i$ of Ω.

Note: Each $\omega \in \Omega$ should be thought of as a specification of all parameters of interest to any player i (including the moves of all other players). $P \in \mathcal{P}^i$ represents player i's knowledge. Notice that knowledge may differ across players!

Denote by $\sigma^i(\omega)$ the action chosen by player i at the state ω. The information structure is such that $\sigma^i(\omega_1) = \sigma^i(\omega_2) \quad \forall \omega_1,\ \omega_2 \in P \in \mathcal{P}^i$, i.e. the player cannot distinguish between elements of P.

We say that a player i is **Bayes rational** at ω if

$$E(u^i(\sigma)|\mathcal{P}^i)(\omega) \geq E(u^i(a^i, \sigma_{-i})|\mathcal{P}^i)(\omega) \quad \forall a^i \in A^i,$$

where for any $\omega \in P \in \mathcal{P}^i,\ E(x|\mathcal{P}^i)(\omega) = E(x|P)$.

Aumann's Theorem:

(1) Given a correlated equilibrium μ, we can find $(\Omega, \rho, \mathcal{P}, \sigma)$ such that each player is Bayes rational.

(2) If each player is Bayes rational at each state of the world, then $\mu \in \Delta(A)$ induced by σ is a correlated equilibrium distribution. That is, if $\mu(a) = \rho(\{w \mid \sigma^i(\omega) = a^i \quad \forall i\}) \quad \forall a \in A$, then μ is a correlated equilibrium.

Proof. (1) Let

$$\begin{aligned} \Omega &= \{a \in A \mid \mu(a) > 0\}, & \rho &= \mu, \\ \mathcal{P}^i &= \{\{a^i\} \times A^{-i} \mid a^i \in A^i\}, & \sigma^i(\{a^i\} \times A^{-i}) &= a^i. \end{aligned}$$

Then,

$$\rho(a^{-i}|a^i) = \frac{\mu(a^i, a^{-i})}{\sum_{\hat{a}^{-i}} \mu(a^i, \hat{a}^{-i})}.$$

Hence, $\forall i,\ \forall a^i \in A^i\ s.t.\ \mu(a^i) > 0,\ \forall b^i \in A^i$,

$$\begin{aligned} &\sum_{a^{-i}} u^i(a)\rho(a^{-i}|a^i) \geq \sum_{a^{-i}} u^i(b^i, a^{-i})\rho(a^{-i}|a^i). \\ \Rightarrow &\sum_{\omega \in P} u^i(\sigma(\omega))\rho(w|P) \geq \sum_{\omega \in P} u^i(b^i, \sigma^{-i}(\omega))\rho(w|P), \quad P = a^i \times A^{-i} \in \mathcal{P}^i \\ &\Rightarrow E(u^i(\sigma)|\mathcal{P}^i)(\omega) \geq E(u^i(b^i, \sigma^{-i})|\mathcal{P}^i)(\omega) \quad \forall b^i \in A^i, \end{aligned}$$

Where $\omega \in P$.

(2) $\forall i$, define $\mathcal{H}^i(a^i) = \{\omega : \sigma^i(\omega) = a^i\}$. This could be a single information set of player i or a union of many. In any case, it is a coarser partition than $\mathcal{P}^i$.

Every player is Bayes rational, so

$$
\begin{aligned}
& E(u^i(\sigma)|\mathcal{P}^i)(\omega) \geq E(u^i(b^i, \sigma^{-i}(\omega))|\mathcal{P}^i)(\omega) \quad \forall i,\ b^i \in A^i,\ w \in \Omega \\
\Leftrightarrow \quad & \sum_{\omega \in P} u^i(\sigma(\omega))\rho(\omega) \geq \sum_{\omega \in P} u^i(b^i, \sigma^{-i}(\omega))\rho(\omega) \quad \forall i,\ b^i \in A^i,\ P \in \mathcal{P}^i \\
\Rightarrow \quad & \sum_{\omega \in \mathcal{H}^i(a^i)} u^i(\sigma(\omega))\rho(\omega) \geq \sum_{\omega \in \mathcal{H}^i(a^i)} u^i(b^i, \sigma^{-i}(\omega))\rho(\omega) \quad \forall i,\ a^i,\ b^i \in A^i \\
\Rightarrow \quad & \sum_{\omega \in \mathcal{H}^i(a^i)} u^i(a^i, \sigma^{-i}(\omega))\rho(\omega) \geq \sum_{\omega \in \mathcal{H}^i(a^i)} u^i(b^i, \sigma^{-i}(\omega))\rho(\omega) \quad \forall i,\ a^i,\ b^i \in A^i \\
\Rightarrow \quad & \sum_{a^{-i}} u^i(a^i, a^{-i})\mu(a) \geq \sum_{a^{-i}} u^i(b^i, a^{-i})\mu(a) \quad \forall i,\ a^i,\ b^i \in A^i.
\end{aligned}
$$

Example:

Consider the following three-player normal form game.

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

A

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

B

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

C

The pure strategy Nash Equilibria of this game are (D, L, A), (U, R, A), (D, L, C), (U, R, C). Notice that B is never played in a Nash Equilibrium. There is, however, a way to make player 3 (who controls the matrices) play B with probability 1 in correlated equilibrium.

To this end, let $\mu(U, L, B) = \mu(D, R, B) = \frac{1}{2}$. It's trivial to see that the definition of correlated equilibrium is satisfied.

Using Aumann's theorem, each player is Bayes rational in a game augmented by the following information structure:

$\Omega =$ action profiles, $\rho = \mu$, $\mathcal{P}_3 = \big\{\{(A, U, L), (A, U, R), (A, D, L), (A, D, R)\}$, $\{(B, U, L), (B, U, R), (B, D, L), (B, D, R)\}$, $\{(C, U, L), (C, U, R), (C, D, L), (C, D, R)\}\big\}$, etc.

1.10 Nash Equilibrium as an Expression of Something Else

In our analysis of correlated equilibrium, we implicitly assumed each state of the world is associated with some partition element $P \in \mathcal{P}^i$ for every player. We also implicitly assumed that each state of the world is associated with a vector of actions $\sigma(\omega)$. Finally, we assumed that each state of the world is associated with some posterior beliefs $\rho(\cdot|P)$, which in the analysis above were derived from a common prior. We now make these assumptions more explicit to study the relationship between knowledge and Nash equilibrium. The analysis in this section of the notes will closely follow Chapter 5 of Osborne and Rubinstein.

Definition: An information function P associates with every state $\omega \in \Omega$ a nonempty subset $P(\omega)$ of Ω.

Definition: An information function is partitional if and only there is a partition $\mathcal{P}$ of Ω such that for any $\omega \in \Omega$ the set $P(\omega)$ is the element of the partition that contains ω.

Note: If an information function is partitional, we can work either with the information function itself or with the associated information partition $\mathcal{P}$, as we did in our analysis of correlated equilibrium.

Note: If an information function is partitional, then the following two conditions are satisfied:

1. $\omega \in P(\omega)$ for every $\omega \in \Omega$;
2. If $\omega' \in P(\omega)$, then $P(\omega') = P(\omega)$.

The first condition simply says that the decision maker is never sure that the state is different from the true state. The second condition says that the decision maker can use the consistency of other states with his information to make inferences. For example, assume that the true state is ω. Also, contrary to the second condition, assume that $\omega' \in P(\omega)$ but there is $\omega'' \in P(\omega')$ such that $\omega'' \notin P(\omega)$. Then, the decision maker can argue that the true state cannot be ω'. If the second condition needs to be satisfied, ω' is ruled out.

Homework: Show if an information function satisfies the two conditions above, then it is partitional.

If $P(\omega) \subset E$ for some set E, we say that the decision maker knows E. This motivates our definition of the knowledge function.

Definition: The decision maker's knowledge function associated with every set E the set of states at which the decision maker knows E:

$$K(E) = \{\omega \in \Omega : P(\omega) \subset E\}.$$

Homework: Chapter 5 of Osborne and Rubinstein states several intuitive conditions that the knowledge function must satisfy. Verify these conditions.

So far, we've been talking about knowledge in the context of a single individual. Let's move to a context with multiple decision makers. We say that an event is "common knowledge" if each individual knows that each individual knows the event, each individual knows that each individual knows that each individual knows the event, and so on. Let's formalize this idea.

Definition: Let K^1 and K^2 be the knowledge functions of individuals 1 and 2 for the set Ω of states. An event E is common knowledge between 1 and 2 at the state ω if $\omega \in K^1(E)$, $\omega \in K^2(E)$, $\omega \in K^1(K^2(E))$, $\omega \in K^2(K^1(E))$, $\omega \in K^1(K^2(K^1(E)))$, $\omega \in K^2(K^1(K^2(E)))$, etc.

Example: Assume $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$. The event $E = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ is not common knowledge at any state $\omega \in \Omega$ given the following information partitions:

$$\mathcal{P}^1 = \{\{\omega_1, \omega_2\}, \{\omega_1, \omega_2, \omega_3\}, \{\omega_6\}\},$$

$$\mathcal{P}^2 = \{\{\omega_1\}, \{\omega_2, \omega_3, \omega_4\}, \{\omega_5\}, \{\omega_6\}\}.$$

Now, let's apply these ideas in the context of normal form games. To this end, let's fix our attention on a normal form game $G = (I, (A^i)_{i=1,\ldots,n}, (u^i)_{i=1,\ldots,n})$. Let Ω be a set of states. Each state is a complete description of everything needed to know by each player to make his decision. Formally, each state $\omega \in \Omega$ specifies the following for every player i:

- $P \in \mathcal{P}^i$, which describes player i's knowledge in state ω
- $\sigma^i(\omega)$, which describes player i's action in state ω
- $\mu^i(\omega)$, which describes the beliefs of player i over the actions A^{-i} of the other players.

In our analysis of correlated equilibrium above, we assumed that the third object is derived from a common prior $\rho(\omega)$. Indeed, Aumann's result is that if beliefs are derived from a common prior, and that if every player is Bayes rational at every state ω, then the implied correlated strategy is a correlated equilibrium! Notice that Aumann's result implicitly assumes common knowledge of rationality because every event that is true at every state is common knowledge. It is perhaps easiest to see this with two players. Let R denote the event that every player is rational. By assumption, $R = \Omega$, so $\Omega = K^1(\Omega) = K^2(\Omega) = K^1(K^2(\Omega)) = K^2(K^1(\Omega))$, and so forth. Hence, any element of Ω is an element of the infinite intersection of the K's. I.e., there is common knowledge of rationality at every state.

We now relax the assumption of common priors. Our question is what "epistemic" (knowledge) conditions characterize Nash equilibrium behavior.

Proposition (Epistemic Foundations of Nash Equilibrium): Assume that in the state $\omega \in \Omega$, each player i (1) knows the other players' actions, (2) has beliefs consistent with this knowledge, and (3) is rational. Then, $(\sigma^i(\omega))_{i=1}^N$ is a Nash equilibrium of G.

Let's formalize the three assumptions above. To know the other players' actions means $P^i(\omega) \subset \{\omega' \in \Omega : \sigma^{-i}(\omega') = \sigma^{-i}(\omega)\}$. I.e., player i only considers the states where others choose the same actions as those selected by $\sigma^i(\omega)$.

To have beliefs consistent with one's knowledge means that support of $\mu^i(\omega)$ is a subset of $\{a^{-i} \in A^{-i} : \sigma^{-i}(\omega') = a^{-i} \text{ for some } \omega' \in P^i(\omega)\}$.

To be rational means $\sigma^i(\omega)$ maximizes player i's utility given the belief $\mu^i(\omega)$.

Proof: By (3), $\sigma^i(\omega)$ maximizes player i's utility given the belief $\mu^i(\omega)$. By (1), $\sigma^{-i}(\omega') = \sigma^{-i}(\omega)$ for all $\omega' \in P^i(\omega)$. By (2), the support of $\mu^i(\omega)$ is the point $\{\sigma^{-i}(\omega)\}$. Hence, $\sigma^i(\omega)$ maximizes player i's utility against $\sigma^{-i}(\omega)$. This is the definition of (pure-strategy) Nash equilibrium.

Some people didn't like the the assumption that everyone knows everyone else's strategy and formulated different epistemic characterizations of Nash equilibrium. This motivates our second proposition.

Proposition 2 (Epistemic Foundations of Nash Equilibrium): Assume that $|I| = 2$ and in the state $\omega \in \Omega$, each player i (1) knows the other player's belief, (2) has beliefs consistent with this knowledge, and (3) knows that the other player is rational. Then, the mixed strategy profile $(s^1, s^2) = (\mu^2(\omega), \mu^1(\omega))$ is a Nash equilibrium of G.

Again, let's formalize the assumptions of this proposition. Knowing the other player's belief means $P^i(\omega) \subset \{\omega' : \mu^j(\omega) = \mu^j(\omega')\}$.

Having beliefs consistent with this knowledge means the support of $\mu^i(\omega)$ is a subset of $\{a^{-i} \in A^{-i} : \sigma^{-i}(\omega') = a^{-i}$ for some $\omega' \in P^i(\omega)\}$.

To know that the other player is rational means to know that for any $\omega' \in P^i(\omega)$, $\sigma^j(\omega')$ maximizes j's utility against $\mu^j(\omega')$.

Proof: It suffices to show that any action in the support of $\mu^1(\omega)$ is a pure best response to $\mu^2(\omega)$. So, take any action a^2 in the support. By (2), $a^2 = \sigma^2(\omega')$ for some $\omega' \in P^1(\omega)$. By (3), $\sigma^2(\omega')$ maximizes 2's utility against $\mu^2(\omega')$. By (1), $\mu^2(\omega') = \mu^2(\omega)$. This finishes the proof.

The following example shows that **Proposition 2** doesn't have an analogue in games with more than two players. In particular, we will consider a three player example in which all three conditions of the proposition are satisfied but the players' beliefs do not constitute a mixed strategy Nash equilibrium.

Example: In the example below, there are three players and six states given by $\Omega = \{\alpha, \beta, \gamma, \delta, \epsilon, \xi\}$. The players' information partitions and actions for each state are specified in the figure below. We assume that posterior beliefs are derived from the common prior given in the first row.

Consider the state $\delta \in \Omega$. Condition (2) of **Proposition 2** is clearly satisfied at this state since beliefs are derived from the common prior. It is left as a homework exercise to show that the other conditions are satisfied, as well. On the other hand, the players' beliefs in this state do not constitute an equilibrium. In fact, they do not even coincide (compare the beliefs of player 1 and player 2 about player 3).

	L	R
U	$2,3,0$	$2,0,0$
D	$0,3,0$	$0,0,0$

A

	L	R
U	$0,0,0$	$0,2,0$
D	$3,0,0$	$3,2,0$

B

State	α	β	γ	$\boldsymbol{\delta}$	ϵ	ξ
Probability×63	32	16	8	**4**	2	1
1's action	U	D	D	$\boldsymbol{D}$	D	D
2's action	L	L	L	$\boldsymbol{L}$	L	L
3's action	A	B	A	$\boldsymbol{B}$	A	B
1's partition	$\{\alpha\}$	$\{\beta$	$\gamma\}$	$\{\boldsymbol{\delta}$	$\epsilon\}$	$\{\xi\}$
2's partition	$\{\alpha$	$\beta\}$	$\{\gamma$	$\boldsymbol{\delta}\}$	$\{\epsilon$	$\xi\}$
3's partition	$\{\alpha\}$	$\{\beta\}$	$\{\gamma\}$	$\{\boldsymbol{\delta}\}$	$\{\epsilon\}$	$\{\xi\}$

2 Dynamic Games

An extensive form (dynamic) game has the following basic ingredients: set of players I (possibly including nature), set of nodes X, set of final nodes Z, and set of utility functions $u^i : Z \to \mathbb{R}$. If nature controls the initial node, we denote this node by o. Any description of an EFG must also specify:

- How the nodes are ordered;
- How the nodes are allocated among players, including nature;
- How nature determines its moves;
- What information the players have;
- What moves the player can make at every node.

Ordering of nodes

Formally, we require $\succsim$ a partial order over the set of nodes ($x \succsim y$ means x comes before y) with the following properties:

1. $\succsim$ is reflexive, antisymmetric, transitive;
2. $\forall x \quad 0 \succsim x$;
3. $\forall x \in X$, $p(x)$=the set of predecessors of x=$\{y : y \succsim x\}$ is linearly ordered by $\succsim$.

Players' partition

A players' partition is partition $(P^1, P^2, ..., P^N)$ of the set of non-terminal nodes:

1. $\forall i \neq j$, $P^i \cap P^j = \emptyset$;
2. $\cup_i P^i = X \setminus \{o, Z\}$.

Nature's move, if nature is a player

The probability of nature's move is a probability distribution $p_o \in \Delta(S(o))$, where $S(o)$ is the set of immediate successors of the initial node.

Information partition

For every player i, the information partition $U^i = \{u, v, w, ...\}$ is a partition of P^i. That is,

1. $u \cup v \cup w \cup ... = P^i$;
2. $u \neq v \Rightarrow u \cap v = \emptyset$.

Players' choices

Let $S(x)$ the set of immediate successors of node x. For any information set u, let $\cup_{x \in u} S(x)$ to be the set of immediate successors of u. For all $i, u \in U^i$, C_u is the set of choices i can make at u. This is a partition of $\cup_{x \in u} S(x)$ such that each set contains one and only one of the elements of $S(x)$ for every $x \in u$.

Additional properties

We might want to impose some additional intuitive properties on extensive form games. In particular, we might want to rule out situations where a player doesn't remember if he moved or not, and situtations where the player has different numbers of actions at different nodes in the same information set.

Formally, we might assume that the game is **linear**. This means that $\forall i,\ \forall u \in U^i,\ \forall z \in Z, \quad \#\{p(z) \cap u\} \leq 1$. In words, the path of predecessors of any terminal node intersects any information set at most once. This rules out situations where players forget whether or not they moved.

We might also assume that $\forall i,\ \forall u \in U^i,\ \forall x_1, x_2 \in u, \quad \#S(x_1) = \#S(x_2)$. This rules out situations where different nodes in the same information set have different numbers of actions. This would be nonintuitive since we expect that a player would be able to distinguish between such nodes.

Finally, we might want to assume that players remember what they did. Formally, a game satisfies **perfect recall** if there exists no player i with information sets u and v and a choice $c \in C_u$ such that $x_1, x_2 \in v$, x_1 follows c but x_2 does not follow c.

2.1 Strategies in Extensive Form Games

Definition (Pure Strategies): A pure strategy specifies a choice at every information set. Formally, the set of pure strategies $A^i = \times_{u \in U^i} C_u = \{(a^i_u)_{u \in U^i} : a^i_u \in C_u \quad \forall u\}$.

Definition (Mixed Strategies): A mixed strategy is a probability distribution on A^i. Formally, $S^i = \Delta(A^i)$.

2.2 Normal Representation

Every extensive form game has a normal form representation, with the same set of players I and $A^i = \times_{u \in U^i} C_u$, as defined above. Every action profile leads to a terminal node, and $u^i(z)$ at that node gives player i's utility. If nature is one of the players, then for every pure strategy profile $(a^1, ..., a^N)$ of the players, there is a unique probability on the final nodes given by $Pr_{p_0,a} \in \Delta(Z)$ and hence an expected utility vector. Thus, in general, the NFG representation of an EFG has:

- $I = \{1, ..., N\}$;
- For every i, A^i the set of pure strategies;
- $u^i(a) = E_{Pr_{p_0,a}} u^i(\cdot)$.

Note: We say that two pure strategies b^i and c^i are *payoff equivalent* if $u^j(b^i, a^{-i}, d^i) = u^j(c^i, a^{-i}) \; \forall a^{-i} \in A^{-i}, \quad \forall j \in I$. I.e., no player cares if i uses b^i or c^i. We can replace each set of payoff-equivalent strategies with a single strategy and relabel. The result is called the (purely) reduced NFG representation.

Example: See class notes.

2.3 NE of EFG

Definition (NE of EFG): The NE of the EFG is the NE of the mixed extension of the associated NFG.

Example: See class notes.

Theorem: The set of NE of any EFG is non-empty.

Proof: Any EFG has a NFG representation. We showed that the set of NE of any NFG is non-empty earlier.

2.4 Behavioral Strategies

Arguably, mixed strategies is not the most natural way of thinking about strategies in an EFG.

Definition (Behavioral Strategies): A behavioral strategy specifies a probability distribution for every information set for every player that the player uses to randomize. Formally, the set of behavioral strategies is of player i is:

$$B^i = \{(b^i_u)_{u \in U^i} : b^i_u \in \Delta(C_u) \quad \forall u \in U^i\}.$$

Examples: See class notes.

Definition (Equivalent Strategies): We say that two strategies are equivalent if they induce the same distributions on final nodes for any pure strategy profile of the other players.

In general, mixed and behavioral strategies are not equivalent (see example from class). They are, however, equivalent in games of perfect recall.

2.5 Kuhn's Theorem

Theorem (Kuhn's Theorem): In any EFG where player i has perfect recall, for any mixed strategy of player i there is an equivalent behavioral strategy and vice versa.

Proof: See Osborne in Rubinstein.

Given a mixed strategy s^i, we define the behavioral strategy as follows:

$$b_u^i(c) = \frac{\text{Probability of reaching } u \text{ and making choice } c \text{ with } s^i}{\text{Probability of reaching } u \text{ with } s^i}.$$

Given a behavioral strategy b^i, we define the mixed strategy as follows:

$$s^i(a) = \Pi_{u \in U^i} b_u^i(c_u(a)),$$

where $c_u(a)$ is is the choice specified by a at information set u.

Example: For a two-player example, see class notes, or the first homework problem of Homework 11.

Definition (NE in behavioral strategies): A Nash Equilibrium in behavioral strategies is a behavioral strategy profile $b = (b^1, b^2, ..., b^N)$ such that for every player i, b^i is a best response to b^{-i}.

Example: See class notes.

Note that as a corollary of Kuhn's Theorem, every EFG with perfect recall has a NE in behavioral strategies. Indeed, the theorem says that it doesn't matter whether we think of behavior in terms of mixed or behavioral strategies if every player has perfect recall. We can write this down as a Theorem:

Theorem: In any EFG of Perfect Recall for every player, for every NE in mixed strategies, there is a NE in behavioral strategies that induces the same probability on final nodes (and same payoff outcomes) and vice versa.

2.6 Subgame Perfect Nash Equilibrium

This equilibrium concept is motivated by the observation that NE sometimes does not make intuitive predictions in extensive form games.

Example: As motivation, consider the following game. There are two players: Firm E (Entrant) and Firm I (Incumbent). Firm E chooses whether to enter a market or not. If it does not enter, the payoffs are 0 to Firm E and 2 to Firm I (Firm I gets all the profit). If Firm E enters, Firm I has two choices: to accommodate or fight. If

Firm I accommodates, it gets a payoff of 1 while firm E gets a payoff of 2. If Firm I fights, Firm E gets a payoff of -3 and Firm I gets a payoff of -1.

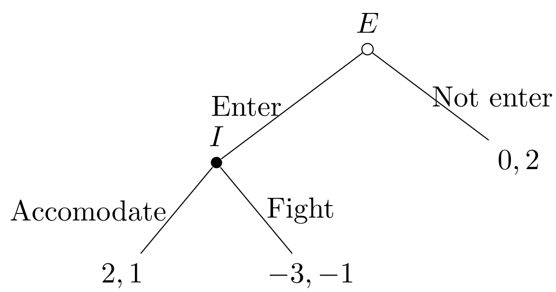

Figure 1: Example: Market entry game

The normal form of this game is given below:

	Accomodate	Fight
Enter	2,1	-3,-1
Not enter	0,2	0,2

The two pure strategy equilibria of this game are (In, Accomodate) and (Out,Fight). While the former equilibrium is reasonable, the latter one is not. This is because Player E knows that if he plays In, there is no way that Player I will decide to fight.

Definition (Subgame): A subgame is the game beginning at every node if it really is a game, i.e. if every player knows the game.

Formally, let $K_x = \{y : x \succsim y\}$. This is the set of nodes following x.

- $\succsim_x$ is the restriction of $\succsim$ to K_x.
- $P^i_x = P^i \cap K_x$.
- $Z_x = Z \cap K_x$.
- $u^i_x(\cdot) = u^i(\cdot)$ defined on Z_x.

A subtree is a subgame if whenever $x_1, x_2 \in u$ for some $u \in U^i$ and $x_1 \in K_x$, it is also the case that $x_2 \in K_x$.

Definition (SPNE): $b \in B$ is a Subgame Perfect Nash Equilibrium (SPNE) is its restriction to the subgame G_x is a Nash equilibrium in G_x for all subgames G_x of G.

Theorem: SPNE exists.

Sketch of proof: To prove the theorem, we can construct a SPNE of any extensive form game G using the following backward induction procedure:

1. Take all the minimal subgames of G (minimal=has no subgames).
2. For each of these subgames, find its behavioral NE strategy.
3. Replace the original game with the game where the minimal subgames are removed and the nodes are assigned expected payoffs from the associated behavioral strategy NE.
4. Repeat.

For details, see the class notes.

2.7 Sequential Equilibrium

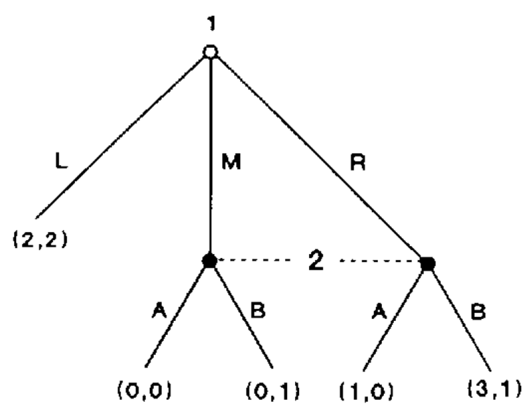

Figure 8.1

Consider the game above. (L,A) and (R,B) are both SPNE (not that the only subgame of this game is the game itself). However, if Player 2 is given the choice to move, choosing B is optimal regardless of what beliefs he has about the nodes in his information set. To deal with this problem, we define a new equilibrium concept which includes beliefs explicitly in the definition. This equilibrium concept will require that, at every information set, every player is behaving optimally given her beliefs.

Definition: A belief system μ is a map $\mu : X \to [0,1]$ s.t. $\sum_{x \in u} \mu(x) = 1$ for all $i, u \in U^i$. **Definition:** An assessment is a tuple (b, μ) where $b \in B$ and μ is a belief system.

Definition: The set of ϵ-fully mixed behavioral strategies of player i is

$$B^i_\epsilon = \{b^i \in B^i : b^i_u(c) \geq \epsilon \quad \forall c \in C_u,\ \forall u \in U^i\}.$$

Definition (Consistent Assessments): An assessment (b, μ) is consistent if there exists a sequence (b_n, μ_n) s.t.

1. b_n is fully mixed $\forall n$;
2. $(b_n, \mu_n) \to (b, \mu)$;
3. $\mu_n(x|u) = \frac{Pr_b(x)}{Pr_b(u)} \quad \forall n, x, u.$

In words, the third condition says that each μ_n is derived from b_n using Bayes' Rule.

Example: Consider the game below.

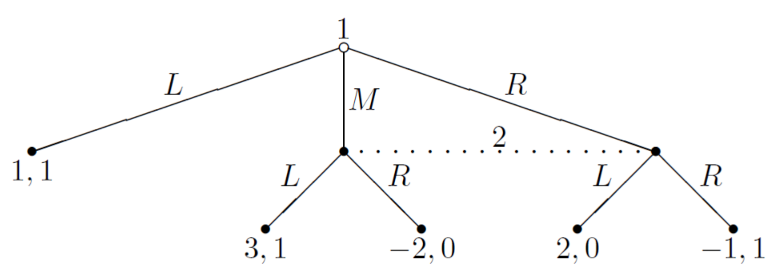

Consider $b = ((1,0,0),(0,1))$ and μ s.t. the distribution over Player 2's nodes is $(\alpha, 1-\alpha)$. The claim is that this assessment is consistent for any $\alpha \in (0,1)$. Let $b^1_n = (1-\epsilon_n, \alpha\epsilon_n, (1-\alpha)\epsilon_n)$ and $b^2 = (\epsilon_n, 1-\epsilon_n)$. Note that $\mu_n(x|u) = \frac{\alpha\epsilon_n}{\epsilon_n} = \alpha$ is the belief of Player 2 about his first node derived using Bayes' Rule.

Definition (Sequential Rationality): An assessment (b, μ) is sequentially rational if

$$E(u^i(b^i, b^{-i}; \mu)|u) \geq E(u^i(d^i, b^{-i}; \mu)|u) \qquad \forall i, u \in U^i, d^i \in B^i.$$

Definition (Sequential Equilibrium): An assessment (b, μ) is a sequential equilibrium if it is consistent and sequentially rational.

Example: Consider the game in the previous example. The assessment (b, μ), where $b = ((1,0,0),(0,1))$ and μ s.t. the distribution over Player 2's nodes is $(\alpha, 1 - \alpha)$, is a sequential equilibrium if $\alpha \leq 1/2$. Question: Can you find other sequential equilibria?

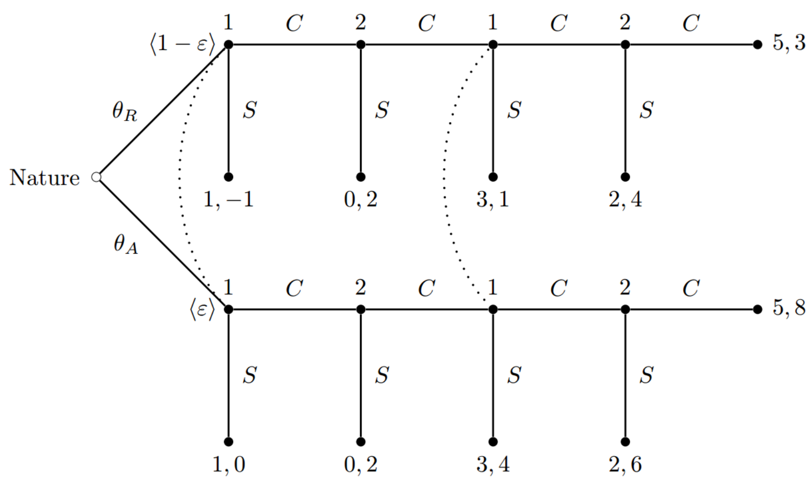

Example: Consider the variant of the centipede game shown above. With probability $\epsilon \in (1/9, 1/3)$, nature makes Player 2 into an altruistic type. Player 1 does not know what kind of Player 2 he is facing, although he knows the probability that Player 2 is altruistic. Find the sequential equilibria of this game.

3 Repeated Games

The big message of the theory of repeated games has two parts. The message is:

- When the game is repeated, you get a lot of "new" Nash Equilibria (a loss of predictive of power),
- When the game is repeated, it's possible to sustain more cooperation than when the game is played once.

Both parts of the message are captured by a set of results that are referred to as the "folk theorems." Let's spend some time trying to understand what these theorems say.

In a repeated game,[1] the fundamental building block is the *stage game*, which is a normal form game $\{I, \{A^i\}, \{g^i\}\}$. This stage game can be played finitely or infinitely many times.

Finitely repeated games

This model is appropriate if the termination time is well-known by all the players.

Infinitely repeated games

This model is appropriate if the termination time is random and not known by the players precisely.

Let $a_t = (a_t^1, ..., a_t^I)$ be the actions played in period t. The game begins in period 0 with null history h_0. For $t \geq 1$, $h_t = (a_0, a_1, ..., a_t)$. Let $H_t = A^t$ be the space of period-t histories.

Definition (pure strategies): A pure strategy of player i specifies a map

$$\sigma_t^i : H^t \to A^i$$

for every period. Thus, a pure strategy is $(\sigma_0^i, \sigma_1^i, \sigma_2^i, ...)$.

Definition (behavioral strategies): A behavioral strategy of player i specifies a map

$$b_t^i : H^t \to S^i$$

[1] These notes are largely based on Chapter 5 of Fudenberg and Tirole.

for every period. Thus, a behavioral strategy is $(b_0^i, b_1^i, b_2^i, ...)$.

Note that a subgame begins at every period of play. We will use this fact later to think about subgame perfect equilibria.

3.1 Infinitely Repeated Games

We focus on the case where players discount future utilities with a discount factor $\delta < 1$.

Given a profile of behavioral strategies b, player i's utility is

$$(1-\delta)E_b[g^i(b_0(h_0)) + \delta g^i(b_1(h_1)) + \delta^2 g^i(b_2(h_2))...] = (1-\delta)E_b\left[\sum_{t=0}^{\infty} \delta^t g_i(b_t(h_t))\right].$$

I.e., it is the expected discounted stream of payoffs with b multiplied by $1-\delta$. The reason for this multiplication is that we want the units to be in terms of the stage game payoffs, i.e. we are in some sense thinking of the average payoffs.

Example Consider an infinitely repeated game where both player have a discount factor $\delta < 1$ and the stage game is the following battle of the sexes game:

	L	R
T	3,1	0,0
B	0,0	1,3

Assume Player 1 chooses B in every period and Player 2 chooses R in every period. Then, Player 1's discounted payoff is:

$$1 + \delta^1 + \delta^2 + \delta^3 + \delta^4 + ... = \frac{1}{1-\delta}.$$

Multiplied by $1-\delta$, we get 1, Player 1's average payoff in the game.

Observation: Let $\hat{\alpha}$ be a Nash Equilibrium of the stage game (a "static equilibrium"). Then, the strategy profile in which every player i plays $\hat{\alpha^i}$ in every period of the repeated game is a Nash Equilibrium of the repeated game. Moreover, if the game has m static equilibria $\{\hat{\alpha_j}\}_{j=1}^m$, then any map from time periods to indices $\{1, ..., m\}$ implies a Nash Equilibrium of the repeated game in which every player plays $\hat{\alpha^i_{j(t)}}$ in period t.

The reason for the observation is: The repeated game strategies described above do not depend on what any player has done in any previous period. Take, for instance, the example where every player i plays $\hat{\alpha^i}$ in every period of the repeated game. When player i finds himself in period t, he asks himself the question: "What is the best thing for me to do now, given that the other players are playing according to $\hat{\alpha^{-i}}$?" Since $\hat{\alpha}$ is a Nash Equilibrium, the answer is to choose $\hat{\alpha^i}$.

What this observation shows is that turning a static game into a repeated game does not make the set of equilibria smaller. In fact, it makes it bigger, as we will now see.

3.2 Folk Theorems for Repeated Games

The folk theorems say that any feasible, individually rational payoffs can be sustained in equilibrium in a repeated game if players are sufficiently patient. To understand what this means, we need to understand the words "feasible" and "individually rational".

Individually rational payoffs

Definition (Minimax value): Define player i's reservation utility or minimax value as

$$\underline{v}^i = \min_{s^{-i}} \left[\max_{s^i} g^i(s^i, s^{-i}) \right].$$

In words, this is the lowest payoff that the other players can hold player i to, assuming that he correctly foresees what they do and plays a best response. Let (m_i^i, m_i^{-i}) be the strategy profile so that $g^i(m_i^i, m_i^{-i}) = \underline{v}^i$.

Example: Find the minimax values of Player 1 and Player 2 in the following game:

	L	R
T	-2,2	1,-2
M	1,-2	-2,2
B	0,1	0,1

Observation: Player i's payoff is at least $\underline{v}^i$ in any static equilibrium and in any Nash Equilibrium of the repeated game.

Proof: In a static equilibrium $\hat{\alpha}$, $\hat{\alpha^i}$ is a best response to $\hat{\alpha^{-i}}$. Thus,

$$g^i(\hat{\alpha^i}, \hat{\alpha^{-i}}) \geq g^i(m_i^i, \hat{\alpha^{-i}}) \geq \min_{s^{-i}} g^i(m_i^i, s^{-i}) = g^i(m_i^i, m_i^{-i}).$$

This shows the first part of the observation.

For the second part of the observation, consider a Nash Equilibrium $\hat{\alpha}$ of the repeated game. One feasible (not necessarily optimal) strategy for player i is to maximize $g^i(s^i, \hat{\alpha_t^{-i}})$ in every period. Now,

$$\max_{s^i} g^i(s^i, \hat{\alpha_t^{-i}}) \geq \min_{s^{-i}} \left[\max_{s^i} g^i(s^i, s^{-i}) \right] = \underline{v}^i.$$

This shows the second part of the observation.

Feasible payoffs

We define the set of feasible payoffs as the convex hull of the payoffs in the stage game:

$$V = \text{convex hull}\{v | \exists a \in A \text{ with } g(a) = v\}.$$

Remember that not all of these payoffs are feasible in the stage game! Some convex combinations of stage game payoffs correspond to correlated equilibrium payoffs, that you cannot get with Nash. In repeated games, though, the story is different. You can get things in the convex hull even without correlated equilibria (assuming that the players are sufficiently patient).

Definition (Feasible, Individually Rational Payoffs): The set of feasible, strictly individually rational payoffs is

$$\{v \in V | v^i > \underline{v}^i \quad \forall i\}$$

Example: In the game above, the set of feasible, individually rational payoffs is given in the following picture:

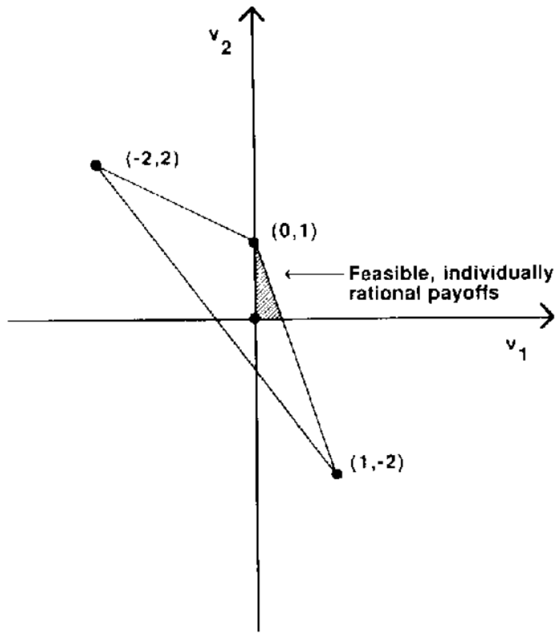

Theorem (Folk Theorem): For every feasible payoff vector with $v^i > \underline{v}^i$ for all players i, there exists $\bar{\delta} < 1$ such that for all $\delta \in (\bar{\delta}, 1)$, there is a Nash Equilibrium $G(\delta)$ of the repeated game with payoffs v.

Proof Assume first that there is a pure strategy profile a s.t. $g(a) = v$. Consider the following strategy for each player i:

- Start by playing a_i;
- Continue playing a_i as long as the previous action was a or the previous action differed from a in two or more components. If in some previous period period player i was the only one not to follow the profile a, then each player j plays m_i^j for the rest of the game.

Can anyone gain by deviating from this strategy profile? Consider the player who deviates in period t, assuming everyone else is being a good boy. In the period the player deviates, he gets AT MOST

$$(1-\delta^t)v^i + \delta^t \left[(1-\delta)\max_a g^i(a) + \delta \underline{v}^i\right] = (1-\delta^t)v^i + \delta^t(1-\delta)\max_a g^i(a) + \delta^{t+1}\underline{v}^i.$$

For a δ close to 1, the deviating player's payoff is clearly less than v^i (since $\underline{v}^i < v^i$). This finishes the proof for the case where v can be attained in pure strategies.

Now consider the case where v cannot be attained in pure strategies. Here, our proof will make use of public randomizations. A public randomization is a publicly observed coin flip w that the players can use to correlate their actions. While the Folk Theorem can also be proved *without* public randomizations, we use them because it makes the proof easier.

Formally, let $\{\omega_1, \omega_2, ...\}$ be a sequence of uniform draws from the interval $[0, 1]$. The history is now $h_t = (a_0, a_1, ..., a_{t-1}, \omega_0, \omega_1, ..., \omega_{t-1})$. As before, a pure strategy is a sequence of maps from histories to actions.

Let v be a feasible, strictly individually rational payoff vector that cannot be attained in pure strategies. Let $a(\omega)$ be a public randomization with expected payoffs v. Consider the following strategy:

- Start by playing according to $a(\omega)$;
- Continue playing according to $a(\omega)$ as long as everyone played according to $a(\omega)$ in the past, or the number of deviating players was two or greater. If in some previous period period player i was the only one not to play according to $a(\omega)$, then each player j plays m_i^j for the rest of the game.

Define $\bar{\delta}$ as follows:

$$(1-\bar{\delta})\max_a g(a) + \bar{\delta}\underline{v_i} = (1-\bar{\delta})\min_a g(a) + \bar{\delta}v_i.$$

For $\delta > \bar{\delta}$, if some player i deviates in some period t he gets at most:

$$(1-\delta^t)v^i + \delta^t\left[(1-\delta)\max_a g^i(a) + \delta\underline{v^i}\right] \leq (1-\delta^t)v^i + \delta^t\left[(1-\delta)\min_a g^i(a) + \delta v^i\right] \leq v^i$$

This finishes the proof.

3.2.1 Example (Repeated Cournot Oligopoly)

Consider a two firm Cournot Oligopoly in which the demand curve is given by $p(q)$ with $p(q) \to 0$ as $q \to \infty$ and each firm faces a per-unit cost c. The Folk Theorem says that for a large enough discount factor the firms can collude to each produce of one half of the monopoly output by threatening each other to produce competitive equilibrium quantities (i.e., minmaxing. Work out the details for the homework).

There is a potential problem here. Imagine now that the costs of the two firms are different: Firm i faces a per-unit cost of c_i. In principle, it's possible that the quantity q_j that minimaxes firm i is such that c_j exceeds the price. This raises the question of whether firm j would want to commit to following the punishment path. More generally, the issue is that the equilibrium strategies described above are not subgame perfect.

Is it possible to sustain feasible, individually rational payoffs using subgame perfect strategies?

Theorem (Nash Threat Folk Theorem): Let α^* be a static Nash Equilibrium of the stage game with payoffs e. For every feasible payoff vector with $v^i > e^i$ for all players i, there exists $\bar{\delta} < 1$ such that for all $\delta \in (\bar{\delta}, 1)$, there is a *Subgame Perfect Nash Equilibrium* $G(\delta)$ of the repeated game with payoffs v.

Proof: The proof of this version of the Folk Theorem looks very similar to the proof we studied above. Let's assume first that there is an action profile a s.t. $g(a) = v$. Consider the strategy in which each player plays a^i as long as the other players do it. If one of the players deviates, all others play α^*. The deviating player gets at most

$$(1-\delta^t)v^i + \delta^t\left[(1-\delta)\max_a g^i(a) + \delta e^i\right].$$

Since $e^i < v^i$, for large enough δ, the deviation payoff is less than v^i. If there is no action profile with $g(a) = v$, we use public randomizations as before (end of proof).

This version of the Folk Theorem is somewhat more reassuring: It tells us that Cournot oligopolists can collude by threatening each other with static Cournot quantities. These are at least static Nash Equilibria, so commitment is not an issue.

3.3 Imperfect Monitoring

We will now consider a class of repeated games with imperfect monitoring. These are games in which *players don't observe each other's actions.* Instead, they observe some noisy (imperfect) information about what the other players are doing. First, consider a repeated game where player i's payoff $r^i(a^i, y)$ in every period depends on his own action a^i and a public signal y. The other player's actions enter through the distribution $\pi_a(y)$ over the public signal. In every period, the player's expected payoff from the action profile a is given by:

$$g^i(a) = \sum_y \pi_a(y) r^i(a^i, y)$$

We will restrict our attention to strategies that are functions of the observed signals. A (public) history is the history of y's observed in previous periods. Thus,

$$h_t = (y_0, y_1, y_2, ...)$$

A (behavioral) strategy of player i specifies a map

$$b_t^i : H^t \to S^i$$

for every period. Thus, a behavioral strategy is $(b_0^i, b_1^i, b_2^i, ...)$, a sequence of maps from histories to mixed actions.

Example (Noisy Prisoner's Dilemma)

Consider an infinitely repeated prisoner's dilemma (like the one we studied before) with the following slight modification. In every period, there is a probability ϵ of each player independently making a mistake. The signal now is the observed action profile, but it depends noisily on the chosen action profile as follows:

$$\pi_{C,C}(C, C) = (1 - \epsilon)^2$$
$$\pi_{C,C}(C, D) = \pi_{C,C}(D, C) = \epsilon(1 - \epsilon)$$
$$\pi_{C,C}(D, D) = \epsilon^2$$

Thus, $H = A$ and h_t is the sequence of *observed* actions, which may be different from the *intended* actions. This is a repeated game with imperfect monitoring.

From now on, we will change the environment slightly. We will assume that the underlying game is some normal form game G. The players' payoffs are given by the payoffs in the normal form game, but these payoffs are unobserved. All that the players observe is the history of the public signals. We will focus on **perfect public equilibria**. In words, these are equilibria in strategies *that only depend on histories of the public signal*. Our definition will also imply subgame perfection:

Definition A profile $b = (b^1, ..., b^I)$ is a **perfect public equilibrium** if

1. Each b_i is a public strategy;

2. For each t, h_t, the strategies yield a Nash equilibrium from that date on.

Let's study what we can achieve in perfect public equilibria in the following simple example:

Example 2

Depending on whether she chose to cooperate (C_{it}) or defect (D_{it}), subject i's stock of points increases by g_{it} in period t according to the following table:

	C	D
C	15, 15	0, 20
D	20, 0	2, 2

Instead of observing her partner's actions, each player sees a public signal that could go up with probability $p(a_t)$ or down with probability $1-p(a_t)$, with $p(a_t)$ determined as in the table below:

	C	D
C	$\frac{3}{4}$	$\frac{1}{2}$
D	$\frac{1}{2}$	$\frac{1}{2}$

Proposition (Bound on Public Equilibrium Payoffs) Let $\gamma(\delta) = \max\{v_1 + v_2 : v \in E(\delta)\}$, where $E(\delta)$ is the set of public equilibrium payoff vectors of the game with discount factor $\delta < 1$. Then, $\gamma(\delta) \leq 20$ for every δ.

An illustration of this proposition is given below:

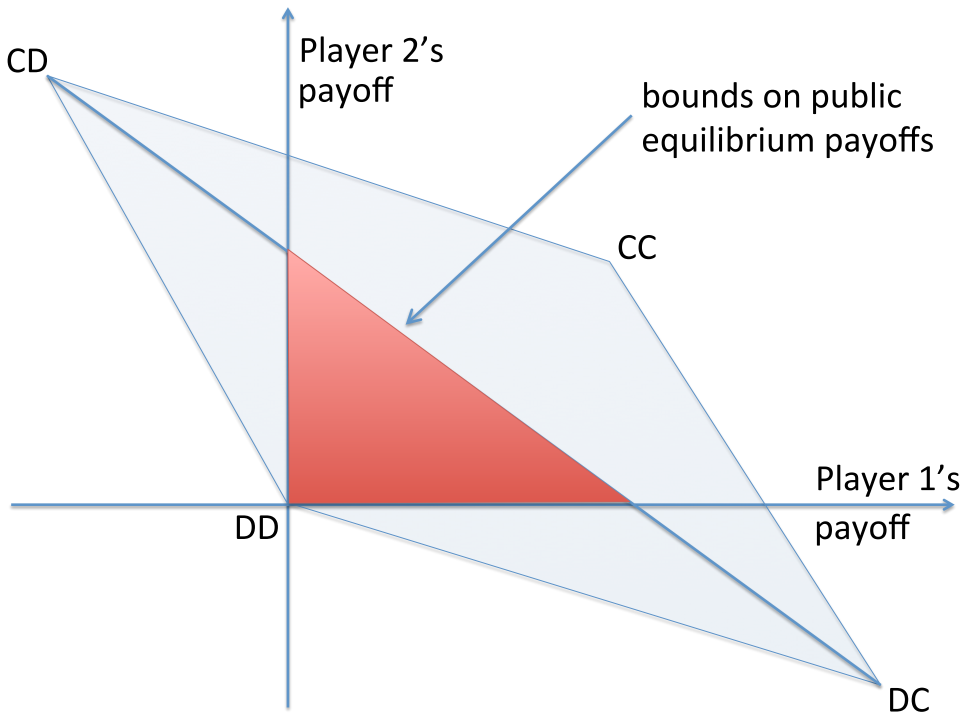

Proof: The proof follows a basic argument by Fudenberg, Levine and Maskin (1994). For a contradiction, assume that $\gamma > 20$. Choose $v \in E(\delta)$ such that $v_1 + v_2 = \gamma$. Player i's utility is given by

$$v_i = (1-\delta)u_i + \delta(pw_i^+ + qw_i^-),$$

where w_i^+ and w_i^- denote the continuation payoffs of player i after a good and a bad signal, respectively, u_i is player i's expected utility today and p and q are, respectively, the probabilities of a good and a bad signal today. Since $\gamma > 20$, it must be the case that the probability of both players cooperating is greater than zero after some history. Let μ_j, where $j \neq i$, denote player j's probability of defection. It will be incentive compatible for player i to cooperate if

$$(1-\delta)(1-\mu_j)15 + \delta\big[\big((1-\mu_j)p_2 + \mu_j p_1\big)w_i^+ + \big((1-\mu_j)q_2 + \mu_j q_1\big)w_i^-\big]$$

$$\geq$$

$$(1-\delta)\big((1-\mu_j)20 + 2\mu_j\big) + \delta\big[\big((1-\mu_j)p_1 + \mu_j p_0\big)w_i^+ + \big((1-\mu_j)q_1 + \mu_j q_0\big)w_i^-\big].$$

Write $\Delta p = (1-\mu_j)(p_2 - p_1) + \mu_j(p_1 - p_0)$. Rearranging yields:

$$(1-\delta)[5(1-\mu_j) + 2\mu_j] \leq \delta \Delta p(w_i^+ - w_i^-),$$

that is, current utility gains from deviating are outweighed future losses in continuation payoffs. Since $p_2 - p_1 = 1/4$ and $p_1 = p_0$, this inequality yields the following upper bound on w_i^-:

$$w_i^- \leq w_i^+ - \frac{1-\delta}{\delta}\frac{5(1-\mu_j)+2\mu_j}{\Delta p} = w_i^+ - \frac{1-\delta}{\delta}\frac{5(1-\mu_j)+2\mu_j}{(1-\mu_j)/4} \leq w_i^+ - 20\frac{1-\delta}{\delta}$$

Substituting into the previous expression for v_i,

$$v_i \leq (1-\delta)u_i + \delta\left[w_i^+ - 20q\frac{1-\delta}{\delta}\right].$$

Therefore,

$$v_1 + v_2 \leq (1-\delta)\left[30 - 40q\right] + \delta\gamma.$$

Since $q \geq 0.25$ and $v_1 + v_2 = \gamma$ by hypothesis, it follows that $\gamma \leq 20$, as claimed. **QED**