ECON 6110

Game Theory Notes

Gabe Sekeres

Spring 2025

Contents

Introduction

This is a graduate-level introduction to game theory and its economic applications. A game is any situation where individuals make choices whose consequences also depend on others' behavior. We will follow the classic Osborne and Rubinstein, and we will follow it fairly directly. Grades will be based on one midterm (30%), five problem sets (four of which count for 20%), and a cumulative final exam (50%). Work in groups for the problem sets, but submit an individual solution. Lecture notes will be posted on Canvas.

1 Static Games

1.1 Strategic Games

A strategic game is a model of interactive decision-making in which each decision-maker chooses a plan of action; choices are made simultaneously (or without knowledge); the plan is chosen to maximize the decision maker's utility, which depends on the action profile. To study these types of situations, we need a formalism.
Definition. A *strategic game* consists of a finite set N of players, for each player i a non-empty set of actions A_i, where the set of all actions is $\mathcal{A} = \bigtimes_{i \in N} A_i$, and for each i a utility function $u_i : A \to \mathbb{R}$. We consider the game as a tuple $\langle N, \{A_i\}, \{u_i\} \rangle$

This is an abstract definition. We can (and sometimes will!) think of games as a map from actions to payoffs, through the *consequence function* $g : \mathcal{A} \to \mathbb{R}^{|N|}$.

The consequences may depend on unknown variables, such as $\omega \in \Omega$. We can either allow a consequence function that depends on ω, $g(a, \omega)$, or by introducing a fictitious player *Nature*.

We can represent simple games (with two players) in matrix form:

	L	R
T	a_1, a_2	b_1, b_2
B	c_1, c_2	d_1, d_2

Here, we have $N = \{r, c\}$, $A_r = \{T, B\}$, $A_c = \{L, R\}$, and payoffs directly of the elements of the matrix.

Some notes:
Remark. Players do not need to choose actions simultaneously, they just need to make decisions independently without knowing the choices of the opponents.
Remark. Rules of the game and preferences are common knowledge; everybody knows, everybody knows that everybody knows, and so on.
Remark. Actions may be fairly complicated contingent plans. For example, r's choice can depend on what c chooses. This expands the above game to the following:

	R, L	R, R	L, R	L, L
T				
B				

Note that this is still a simultaneous (strategic) game! Because the column player chooses without knowing what the other will choose.
Definition. A *solution concept* is a rule that assigns to each game $\langle N, A_i, u_i \rangle$ a prediction of an action profile. We can interpret this normatively (how the game *should* be played) or positively (how the game *is* played). We want to design a solution concept with appealing properties.

Definition. A *(pure) Nash equilibrium* of a strategic game $\langle N, A_i, u_i \rangle$ is a profile $a^\star = (a_1^\star, \ldots, a_n^\star) \in A$ of actions such that for every $i \in N$:

$$u_i(a_i^\star, a_{-i}^\star) \geq u_i(a_i, a_{-i}^\star)$$

for all $a_i \in A_i$.

Essentially, no player has a strictly profitable deviation. This is a stability concept – *unilateral stability* – and it is the weakest stability concept that we have.

An extremely relevant question: how do we get to this? When there is only one, it's very simple (well, PPAD complete, but simple strategically). With multiple, more complicated.

Definition. The *best response correspondence* is

$$B_i(a_{-i}) \coloneqq \{a_i \in A_i : u_i(a_{-i}, a_i) \geq u_i(a_{-i}, a_i') \ \forall \ a_i' \in A_i\}$$

Proposition 1.1. *A Nash Equilibrium is a fixed point of B:*

$$a^\star \text{ is a Nash equilibrium } \iff a_i^\star \in B(a_{-i}^\star) \ \forall \ i$$

Example. *Prisoner's Dilemma*

	Don't Confess	Confess
Don't Confess	$(3, 3)$	$(0, 5)$
Confess	$(5, 0)$	$(1, 1)$

This game has a unique equilibrium! Confess, Confess.

Example. *Dove-Hawk*

	Dove	Hawk
Dove	$(3, 3)$	$(1, 4)$
Hawk	$(4, 1)$	$(0, 0)$

This game has two equilibria! Dove, Hawk and Hawk, Dove. This is fairly problematic, though! It's not so obvious how we should converge to the specific equilibrium – we don't know ex-ante.

Example. *Matching Pennies*

	Heads	Tails
Heads	$(1, -1)$	$(-1, 1)$
Tails	$(-1, 1)$	$(1, -1)$

This game has no pure strategy equilibria! Only equilibria in mixed strategies.

Example. *Cournot Competition* This is more of an application, but actually Cournot stated his model before a Nash equilibrium was formalized as a general tool. Two firms, 1 and 2, simultaneously choose output levels $q_i \in [0, \infty)$. The price is $p(q_1, q_2)$, assumed differentiable. Profit is:

$$u_i(q_1, q_2) = q_i \cdot p(q_1, q_2) - c_i(q_i)$$

For simplicity, let's assume linear demand and costs so $p(q_1, q_2) = \max\{0, 1 - q_1 - q_2\}$, so profit for each firm is

$$\pi_1(q_1, q_2) = q_1(1 - q_1 - q_2) - cq_1$$

$$\pi_2(q_1, q_2) = q_2(1 - q_1 - q_2) - cq_2$$

We can find the reaction function by taking first order conditions, which gets us

$$q_1(q_2) = \frac{1 - c - q_2}{2} \qquad \text{and} \qquad q_2(q_1) = \frac{1 - c - q_1}{2}$$

and setting these equal, we get exactly the levels from Cournot competition canonically! Very nice!

Remark. Given a profile a_{-i}, a best response for i is a set $B_i(a_{-i})$. For a Nash equilibrium, we require a

to choose some $a_i^\star \in B_i(a_{-i}^\star)$. But why? For i it would be equally rational to select any distribution with positive probability in $B_i(a_{-i}^\star)$. Why should we allow this? What are its implications?

Definition. Let's define this formally. Denote by $\Delta(A_i)$ as the set of probability distributions over A_i. An element $\alpha_i \in \Delta(A_i)$ is denoted a *mixed strategy* of player i. A degenerate element of $\Delta(A_i)$ that puts probability 1 on an element is called a pure strategy of i. The expected utility of a player given a profile $\alpha = (\alpha_1, \ldots, \alpha_n)$ is the expected utility:

$$U_i(\alpha) = \sum_{a \in A} \left[\prod_{j \in N} \alpha_j(a_j) \right] u_i(a)$$

We define the *mixed extension* of a strategic game as a tuple $\langle N, \Delta\mathcal{A}, U_i(\alpha) \rangle$. A *mixed strategy equilibrium* of the primal game is a pure strategy equilibrium of the mixed extension.

Example. *Improvised* ("I have no idea what I'm doing" – Marco) We have a game

	$\alpha \cdot R$	$(1-\alpha) \cdot L$
U	$(3,3)$	$(1,4)$
D	$(4,1)$	$(0,0)$

The expected utility of choosing U is $U_1(U, \alpha) = 3\alpha + (1-\alpha)$, and the expected utility of D is $U_1(D, \alpha) = 4\alpha$. They will choose U whenever $\alpha < \frac{1}{2}$, they will choose D whenever $\alpha > \frac{1}{2}$. Whenever $\alpha = \frac{1}{2}$, they are indifferent and will choose any mix! Since the game is symmetric, the same is true for the other player, choosing R and L. We see that there is actually a third equilibrium in this game, where both players mix with probability $\frac{1}{2}$.

Question. Why are we introducing mixed equilibria? Recall that a pure strategy Nash equilibrium may fail to exist. Under plausible assumptions, a mixed strategy equilibrium always exists. To prove this, we can use Kakutani's Fixed Point Theorem.

Theorem 1.1. (Kakutani's Fixed Point Theorem) *Let X be a compact convex subset of $\mathbb{R}^n$ and let $f : X \rightrightarrows X$ be a correspondence, where $f(x)$ is nonempty and convex for all $x \in X$, and the graph of f is closed: for all sequences $\{x_n\}$, $\{y_n\}$ for which $y_n \in f(x_n)$ for all n, $x_n \to x$, and $y_n \to y$, we have $y \in f(x)$. Then there exists $x^\star$ such that $x^\star \in f(x^\star)$.*

So we can use Kakutani's to check for a fixed point of the best-response correspondence, which will of course be a Nash equilibrium. We start with a preliminary result:

Theorem 1.2. *A strategic game $\langle N, \{A_i\}, \{u_i\} \rangle$ has a Nash equilibrium if for all i: the set A_i of actions is a nonempty, compact, and convex subset of a Euclidean space; and the utility u_i is continuous and quasi-concave on A_i.*

Proof. We have that $B_i(a_{-i})$ is nonempty since the expected utility is continuous and A_i is nonempty and compact; we have that the set $B_i(a_{-i})$ is convex since u_i is quasi-concave; and we have that B has a closed graph since $\{u_i\}$ are continuous. Conclusion follows from Kakutani immediately. □

Theorem 1.3. (Nash) *Every finite strategic game has a mixed-strategy Nash equilibrium*

Proof. Consider a finite game (with N players and an action space $A = \bigtimes_{i=1}^N A_i$). Each player's set of mixed strategies is $\Delta(A_i)$. Define the best response correspondence as

$$B_i(\alpha_{-i}) = \underset{\alpha_i \in \Delta(A_i)}{\operatorname{argmax}}\, U_i(\alpha_i, \alpha_{-i})$$

We can further define

$$B(\alpha) = B_1(\alpha_{-i}) \times \cdots \times B_N(\alpha_{-N})$$

and note that

$$B : \Delta(A_1) \times \cdots \times \Delta(A_N) \rightrightarrows \Delta(A_1) \times \cdots \times \Delta(A_N)$$

Any fixed point of B is clearly a Nash equilibrium. It suffices to show that the conditions of Kakutani's Fixed Point Theorem hold. We will take them in turn.

1. Let $|A_i| = k < \infty$. Then we have that $\Delta(A_i)$ is just the probability simplex of dimension $k-1$ (in $\mathbb{R}^k$), and simplices are compact, convex, and nonempty. Since our object of interest is the finite Cartesian product of simplices, B is defined over a compact, convex, and nonempty set mapped to itself.

2. For each i, we have that the expected value of a strategy α_i (holding α_{-i} fixed) is

$$U_i(\alpha_i, \alpha_{-i}) = \sum_{a_i} U_i(a_i, \alpha_{-i})\alpha_i(a_i)$$

which is linear in α_i and thus continuous, so by Weierstrass Theorem it attains a maximum over $\Delta(A_i)$, meaning that $B_i(\alpha_{-i})$ are well-defined and so $B(\alpha)$ is nonempty for all α.

3. For any i and α_{-i}, let β_i and β_i' be any two probability distributions over i's pure strategies such that $\beta_i, \beta_i' \in B_i(\alpha_{-i})$. By definition, β_i and β_i' are both best responses by i to α_{-i}, so i is indifferent between all pure strategies in the set

$$S = \text{supp}(\beta_i) \cup \text{supp}(\beta_i')$$

and thus indifferent to all mixed strategies with support S. Since for any $\theta \in (0,1)$ the mixed strategy $\theta\beta_i + (1-\theta)\beta_i'$ has support S, we have that i is also indifferent between it and β_i and β_i', so $\theta\beta_1 + (1-\theta)\beta_i' \in B_i(\alpha_{-i})$, and so B_i is convex over its domain, meaning that B is convex over its domain.

4. Define sequences $(\alpha_i^t, \alpha_{-i}^t) \to (\alpha_i, \alpha_{-i})$ with $\alpha_i^t \in B_i(\alpha_{-i}^t)$ for all t. Suppose towards a contradiction that $\alpha_i \notin B_i(\alpha_{-i})$, meaning that $\exists\ \tilde{\alpha}_i$ and $\varepsilon > 0$ such that

$$U_i(\tilde{\alpha}_i, \alpha_{-i}) \geq U_i(\alpha_i, \alpha_{-i}) + \varepsilon$$

Then, we have that for sufficiently large t,

$$\begin{aligned} U_i(\tilde{\alpha}_i, \alpha_{-i}) &\geq U_i(\tilde{\alpha}_i, \alpha_{-i}^t) - \frac{\varepsilon}{2} \\ &\geq U_i(\alpha_i, \alpha_{-i}) + \frac{\varepsilon}{2} \\ &> U_i(\alpha_i^t, \alpha_{-i}^t) \end{aligned}$$

where the inequalities follow from the fact that $(\alpha_i^t, \alpha_{-i}^t) \to (\alpha_i, \alpha_{-i})$ and that U is continuous. This contradicts the fact that $\alpha_i^t \in B_i(\alpha_{-i}^t)$, so $B_i(\alpha_{-i})$ (and, by implication $B(\alpha)$) has a closed graph.

Thus, Kakutani's Fixed Point Theorem applies and any finite game has a Nash equilibrium. □

1.2 Bayesian Games

We are often interested in interactions in which there may be some uncertainty about the characteristics of the other players (or the state of nature!). To this end, we model the players' uncertainty by introducing a set Ω of states of nature. *States of nature* are descriptions of a player's relevant characteristics.

Definition. A *Bayesian Game* consists of a finite set N of players, a finite set Ω of states of nature (not necessarily finite later, for simplicity here), and for each player we have a set A_i of actions, a finite set of types T_i and a signal function $\tau_i : \Omega \to T_i$, a probability measure p_i over Ω with $p_i(\tau_i^{-1}(t_i)) > 0\ \forall\ t_i \in T_i$ (the prior belief), and a preference relation $\succeq_i$ over $A \times \Omega$. We have the tuple: $\langle N, \Omega, \{A_i, T_i, \tau_i, p_i, \succeq_i\}_{i \in N}\rangle$.

Often a Bayesian game is presented directly in terms of *types*, and sometimes described in terms of Ω and a signal structure expressed as a conditional distribution over types T_i.

Remark. In this definition, we allow for heterogenous priors. Often we will assume a common prior over Ω.
Definition. A *Nash equilibrium of a Bayesian game* is a Nash equilibrium of the strategic game described as follows: The set of players is the set of pairs (i, t_i) for each $i \in N$ and $t_i \in T_i$. The set of actions of player (i, t_i) is A_i, and the preferences $\succeq_{(i,t_i)}$ are such that

$$a^\star \succeq_{(i,t_i)} b^\star \iff L_i(a^\star, t_i) \succeq_i L_i(b^\star, t_i)$$

where $L_i(a, t_i)$ is a lottery over $A \times \Omega$ that assigns probability $\frac{p_i(\omega)}{p_i(\tau_i^{-1}(t_i))}$ to $(\{a^\star(j, \tau_j(\omega)\}_{j \in N}, \omega)$ if $\omega \in \tau_i^{-1}(t_i)$ and 0 otherwise.
Example. *The Volunteers' Dilemma* We have $N = \{1, \ldots, n\}$ and $T_i = c_i = [0, 1]$, $c_i \sim F(\cdot)$, so we have $f_{-i}(c_{-i}) = \mathbb{P}_{l \neq i} F(c_l)$. Preferences are

$$U_i(a) = \begin{cases} v - c & \text{if } i \text{ volunteers} \\ v & \text{if someone, not } i\text{, volunteers} \\ 0 & \text{nobody volunteers} \end{cases}$$

In our definitions, we have that $\Omega = [0, 1]^n$, $\tau_i : c \to c_i$, and $p_{-i} = F_{-i}(c_{-i})$. The expected utility of choosing to volunteer (V) or not volunteer (NV) is

$$EU_i(a; V) = v - c = [1 - P_{-i}]v + P_{-i}v - c$$
$$EU_i(b; NV) = P_{-i}v$$

So i volunteers if $c \leq [1 - P_{-i}]v = c^\star$. This implies that an agent volunteers with probability $\sigma = F(c^\star)$, and we have that $1 - P_{-i} = v[1 - F(c^\star)]^{n-1}$.
Remark. Bayesian games can also describe situations where there is uncertainty about what other players know.
Example. *Bayesian Game with Uncertainty* Consider a game with $N = \{1, 2\}$ and states $\omega_1, \omega_2, \omega_3$. For Player 1, $\tau_1(\omega_1) = \tau_1(\omega_2) = t_1'$, $\tau_1(\omega_3) = t_1''$, and we have that $(b, \omega_j) \succ_1 (c, \omega_j)$ for $j = 1, 2$, but $(c, \omega_3) \succ_1 (b, \omega_3)$. For Player 2, we have that $\tau_2(\omega_1) = t_2'$ and that $\tau_2(\omega_2) = \tau_2(\omega_3) = t_2''$. Here in state ω_1, 2 knows that 1 strictly prefers b to c, but in state ω_2 2 doesn't distinguish between ω_2 and ω_3, so 2 does not know whether $(b, \omega_j) \succ_1 (c, \omega_j)$ or $(c, \omega_j) \succ_1 (b, \omega_j)$. However, in state ω_1, 1 does not know if 2 knows this fact, because 1 cannot distinguish between ω_1 and ω_2.

1.3 Mixed Equilibria

Example. *Battle of the Sexes* We have the game

	Theater	Music
Theater	$(2, 1)$	$(0, 0)$
Music	$(0, 0)$	$(1, 2)$

This game has two pure strategy equilibria: (T, T) and (M, M). What about mixed strategy equilibria? Assume that 2 chooses T with probability $\alpha_2(T)$ (call it α_2). Then 1's expected utility of choosing T and M respectively is

$$U_1(T, \alpha_2) = 2\alpha_2 + 0 \cdot (1 - \alpha_2) = 2\alpha_2$$
$$U_1(M, \alpha_2) = 0 \cdot \alpha_2 + 1 \cdot (1 - \alpha_2) = 1 - 1\alpha_2$$

so Player 1 prefers T if $2\alpha_2 \geq 1 - \alpha_2 \Longrightarrow \alpha_2 \geq \frac{1}{3}$ and strictly prefers M otherwise. Similarly, assume that 1 chooses T with probability α_1. Then 2's utility of T and M respectively is

$$U_2(T, \alpha_1) = 1 \cdot \alpha_1 + 0 \cdot (1 - \alpha_1) = \alpha_1$$
$$U_2(M, \alpha_1) = 0 \cdot \alpha_1 + 2 \cdot (1 - \alpha_1) = 2 - 2\alpha_1$$

so Player 2 prefers T if $\alpha_1 \geq 2 - \alpha_1 \Longrightarrow \alpha_1 \geq \frac{2}{3}$ and strictly prefers M otherwise. This can be seen graphically as:

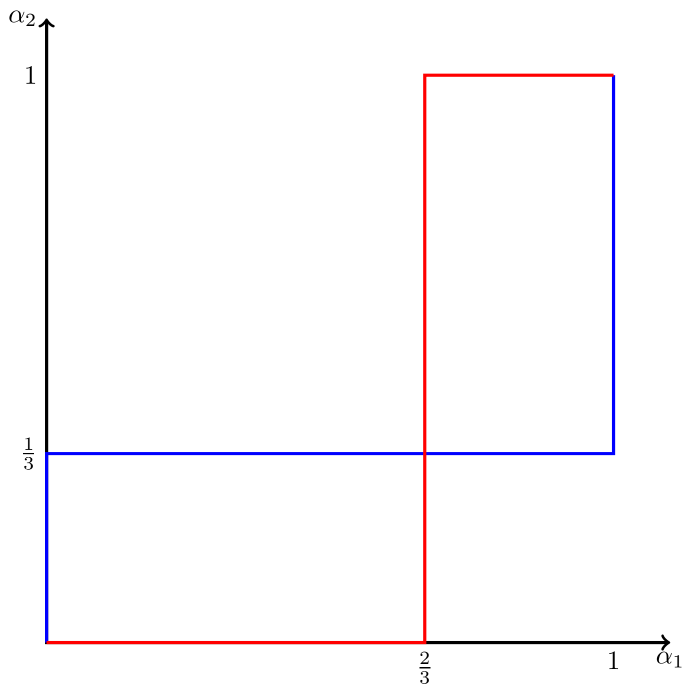

So now the Nash equilibrium is $(2/3, 1/3)$ and the probability that they go to the theater is $\frac{2}{9}$.
Question. What happens if we increase the payoff of player 1 for the theater?

The game is now:

	Theater	Music
Theater	$(5, 1)$	$(0, 0)$
Music	$(0, 0)$	$(2, 1)$

Player 1's expected utility of choosing T and M respectively is

$$U_1(T, \alpha_2) = 5\alpha_2 + 0 \cdot (1 - \alpha_2) = 5\alpha_2$$
$$U_1(M, \alpha_2) = 0 \cdot \alpha_2 + 1 \cdot (1 - \alpha_2) = 1 - 1\alpha_2$$

so Player 1 prefers T if $5\alpha_2 \geq 1 - \alpha_2 \Longrightarrow \alpha_2 \geq \frac{1}{6}$ and strictly prefers M otherwise. Player 2's expected utility of choosing T and M respectively is

$$U_2(T, \alpha_1) = 1\alpha_1 + 0 \cdot (1 - \alpha_1) = \alpha_1$$
$$U_2(M, \alpha_1) = 0 \cdot \alpha_1 + 2 \cdot (1 - \alpha_1) = 2 - 2\alpha_1$$

so Player 2 prefers T if $\alpha_1 \geq 2 - 2\alpha_1 \Longrightarrow \alpha_1 \geq \frac{2}{3}$ and strictly prefers M otherwise. The Nash equilibrium is

now $(2/3, 1/6)$ and the probability that they go to the theater is lower than before!
Example. *Rock, Paper, Scissors* The classic game is represented as:

	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors	$(-1,1)$	$(1,-1)$	$(0,0)$

Let's find all the equilibria! It's clear that there are no pure strategy equilibria. It's also not possible that anyone plays an action with probability 1. Either someone will mix between two actions, or all equilibria will be totally mixed. Let's eliminate the former. Assume WLOG that Player 1 mixes between R and S. Then Player 2 can choose R and guarantee positive payoff, meaning that (since this is zero-sum) Player 2 is guaranteed negative payoff. This is a contradiction, since 1 would then do better by choosing P with probability 1.

Thus, we must have totally mixed strategies. Let's show that the mixed equilibrium is unique. Say that 2 plays $(\sigma_R, \sigma_P, \sigma_S)$, where $\sigma_S = 1 - \sigma_R - \sigma_P$. For 1 to mix, we must have that all of the following are equal:

$$\begin{aligned} u_R &= -\sigma_P + (1 - \sigma_R - \sigma_P) \\ u_P &= \sigma_R - (1 - \sigma_R - \sigma_P) \\ u_S &= -\sigma_R + \sigma_P \end{aligned}$$

Since this is two variables and two independent equations, we generically have a unique solution. Specifically, the first two together imply that $\sigma_R + \sigma_P = \frac{2}{3}$, and the first and third together imply that $1 - \sigma_R = 2\sigma_P$, so $\sigma_R = \sigma_P = \sigma_S = \frac{1}{3}$.

Note that you can use this same strategy all the time, even with asymmetric payoffs.

1.4 Correlated Equilibria

Let's go back to the Battle of the ~~Sexes~~ Friends! We have seen that there are two pure equilibria and a mixed equilibrium. There are other outcomes that can be rationalized. Assume that it rains with probability 1/2, and the players agree to coordinate on T if it rains and M if it doesn't. That would lead to equilibria on the diagonal, and attain higher payoffs in expectation than the mixed equilibrium.

We can make this much more complicated: Suppose there are states $\{x, y, z\}$ with probability 0.4,0.2,0.4 respectively. Suppose that 1 observes $\{x\}$ or $\{y, z\}$, and 2 observes $\{x, y\}$ or $\{z\}$. Assume that 1 believes that 2 plays T if $\{x, y\}$ and M if $\{z\}$, and 2 believes that 1 plays T if $\{x\}$ and M if $\{y, z\}$. This is optimal for 1 if

$$U_1(T : \{x\}) = 2 \geq 0 = U_1(M : \{x\})$$

and

$$U_1(T : \{y, z\}) = \frac{2}{3} \leq \frac{2}{3} = U_1(M : \{y, z\})$$

So this is optimal for 1! Symmetrically, it is also optimal for 2. The probability of (T, T) is now $0.4 > \frac{1}{3}$.
Definition. A *correlated equilibrium* of a strategic game $\langle N, \{A_i\}, \{u_i\}\rangle$ consists of a finite probability space (Ω, π); for each $i \in N$ a partition $\tilde{P}_i$ of Ω (player i's information partition); and for each i a function $\sigma_i : \Sigma \to A_i$ with $\sigma_i(\omega) = \sigma_i(\omega')$ if $\omega, \omega' \in P_i$ for some $P_i \in \tilde{P}_i$ such that for every i and every function $\xi_i : \Omega \to A_i$ with $\xi_i(\omega) = \xi_i(\omega')$ if $\omega, \omega' \in P_i$ for some $P_i \in \tilde{P}_i$ we have:

$$\sum_{\omega \in \Omega} \pi(\omega) \cdot u_i(\sigma_{-i}(\omega), \sigma_i(\omega)) \geq \sum_{\omega \in \Omega} \pi(\omega) \cdot u_i(\sigma_{-i}(\omega), \xi_i(\omega))$$

Remark. Note that the probability space and the partitions are endogenous, part of the equilibrium definition. A Nash equilibrium is a correlated equilibrium, but the opposite is not true, so the set of correlated equilibria is larger than the set of mixed equilibria. Any convex combination of correlated equilibrium payoffs profile is a correlated equilibrium payoff profile of some correlated equilibrium.

Idea: first run a public randomization that identifies which equilibrium to play, and then play that equilibrium.

We may not know which correlation devices are available to the players. Studying the set of correlated equilibria can give us a sense of what outcome we might expect, and what outcome we should not expect. In general, we can assume without loss of generality that the state space coincides with the action space.
Theorem 1.4. Correlated Equilibrium Theorem *Let $G = \langle N, \{A_i\}, \{u_i\} \rangle$ be a finite strategic game. Every probability distribution over outcomes that can be obtained in a correlated equilibrium of G can be obtained in a correlated equilibrium in which: the set of states is A, and for each $i \in N$ player i's information partition consists of all sets of the form $\{a : a_i = b_i\}$ for some $b_i \in A_i$.*

1.5 Evolutionary Equilibrium

Definition. An *evolutionary equilibrium* is a variant of Nash equilibrium concept that has been used to study the evolution of organisms (or other entities). An organism has a possible range of actions B, and is programmed to choose an action $b \in B$. Organisms are paired in anonymous ways to play a game. If an organism chooses b and faces distribution β, the utility is the expected value of $u(b, b')$ where $b' \sim \beta$. As in a two-player symmetric game, $u_1(b, b') = u(b, b')$ and $u_2(b, b') = u(b', b)$. The utility reached in expectation determines the fitness of a type b.
Question. Which steady state should we expect?

The structure of equilibrium here is designed to capture the steady state here, where all organisms take the equilibrium action and no invader can come in with another action.

Intuitively:
Definition. For $b^\star$ to be an *Evolutionary Stable Strategy (ESS)*, we require that

$$(1-\varepsilon)u(b, b^\star) + \varepsilon u(b, b) < (1-\varepsilon)(u(b^\star, b^\star) + \varepsilon u(b^\star, b)$$

for all ε sufficiently small and any $b \in B \setminus \{b^\star\}$.
Remark. The left hand side is the expected utility of a deviation to b, the right hand side is the expected utility of a non-deviator. If this is not satisfied, b has a superior fit to $b^\star$.

Formally:
Definition. A $b^\star$ is said to be an *Evolutionary Stable Strategy (ESS)* if and only if:

1. $b^\star$ is a Nash equilibrium of the symmetric game $\langle \{1, 2\}, (B, B), \{u_i\} \rangle$ with $u_1(a, b) = u_2(b, a)$; and
2. For all $b \neq b^\star$, either $u(b, b^\star) < u(b^\star, b^\star)$ (*i.e.* the equilibrium is strict) or $u(b, b^\star) = u(b^\star, b^\star)$ and $u(b, b) < u(b^\star, b)$

Note that B may be the set of mixed strategies.
Remark. A game may have no ESS. Consider:

	c	d
c	$(1, 1)$	$(1, 1)$
d	$(1, 1)$	$(1, 1)$

However! Every two-player symmetric strategic game in which each player with $|A_i| = 2$ and generic payoffs has an ESS. In general, we have

	c	d
c	(w,w)	(x,y)
d	(y,x)	(z,z)

WLOG, we assume $w < y$ and $z < x$, otherwise we will have a strict pure strategy equilibrium. We must therefore have a fully mixed equilibrium, where

$$w\alpha_c + x(1-\alpha_c) = y\alpha_c + z(1-\alpha_c) \iff \alpha_c = \frac{(z-x)}{(w-y+z-x)}$$

To verify that this is an ESS, we need to show that for any α, $u(\alpha,\alpha) - u(\alpha_c,\alpha) < 0$. This implies that

$$\begin{aligned}
0 &> (\alpha-\alpha_c)[\alpha w + (1-\alpha)x] - (\alpha-\alpha_c)[\alpha_y + (1-\alpha)z] \\
&= (\alpha-\alpha_c)[\alpha(w-y+z-x) + (x-z)] \\
&= (\alpha-\alpha_c)(w-y+z-x)\left[\alpha - \frac{z-x}{w-y+z-x}\right] \\
&= \underbrace{(\alpha-\alpha_c)^2}_{>0}\underbrace{(w-y+z-x)}_{<0}
\end{aligned}$$

1.6 Rationalizability

Motivation. In a Nash equilibrium, we assume that each player optimally responds given their beliefs, and we assume that those beliefs are correct. Each player *knows* the other players' equilibrium behavior. This is absurd, especially in one-shot requirements. We could give up on 'correctness' and rely only on rationality – players actions are optimal based on their beliefs, each player believes that the actions of the other players is a best response to some belief, and this in turn is backed by optimal behavior supported by 'backed' beliefs.
Remark. In most games, rationalizability does nothing. In the Prisoner's Dilemma and Hawk-Dove, the only rationalizable strategies are the Nash strategies.

However, there are situations where it matters. Consider:
Example. *A* 4×4 *Game* The game is:

	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	$0,0$	$0,-2$	$0,0$	$10,-1$

We can immediately see that b_4 is not rationalizable, and from there a_4 is not rationalizable. The game becomes:

	b_1	b_2	b_3
a_1	$0,7$	$2,5$	$7,0$
a_2	$5,2$	$3,3$	$5,2$
a_3	$7,0$	$2,5$	$0,7$

a_1 is a best response if b_3 is played, which is a best response to a_3, which is a best response to b_1, which is a best response to a_1. We have a loop! All of these actions are rationalizable. Additionally, (a_2, b_2) is a Nash equilibrium, so all of the remaining strategies are rationalizable.

We have two equivalent definitions:
Definition. An action $a_1 \in A_1$ is *rationalizable* in the strategic game $\langle N, \{A_i\}, \{u_i\}\rangle$ if there exists:

1. A collection $\{\{X_j^t\}_{j\in N}\}_{t=1}^{\infty}$ of sets with $X_j^1 \subseteq A_j$ for all j and t;
2. A belief μ_i^1 of player i whose support is a subset of X_{-i}^1; and
3. For each player $j \in N$ and $t \geq 1$ and each $a_j \in X_j^t$ a belief $\mu_j^{t+1}(a_j)$ of player j with support X_{-j}^{t+1};

such that

- $a_i = a_i^0$ is a best response to the belief μ_i^1 of player i, so for every j and $t \geq 1$, every action $a_j \in X_j^t$ is a best response to the belief $\mu_j^{t+1}(a_j)$ of player j
- The sets X_j^1 for $j \in N \setminus \{i\}$ are defined as the set of a'_j such that there is an a_{-i} in the support of $\mu_i^1(a_1)$ for which $a'_j = \{a_{-i}\}_j$, *i.e.* a'_k is the jth element of a_{-i} (and $X_i^1 \neq \varnothing$ by convention)
- The sets X_j^t for $t \geq 2$ are defined as the set of a'_j such that there is some player $k \in N \setminus \{j\}$ some action $a_k \in X_k^{t-1}$ and some a_{-k} in the support of $\mu_j^t(a_k)$ for which $a'_j = \{a_{-k}\}_j$.

This definition is a mess! The following is much easier to remember and check:
Definition. An action $a_i \in A_i$ is *rationalizable* in the strategic game $\langle N, \{A_i\}, \{u_i\}\rangle$ if for each $j \in N$ there is a set $Z_j \subseteq A_j$ such that:

1. $a_i \in Z_i$
2. Every action $a_j \in Z_j$ is a best response to a belief $\mu_j(a_j)$ of player j whose support is a subset of Z_{-j}.

Proposition 1.2. *These two definitions are equivalent,*

Proof. If $a_i \in A_i$ is rationalizable according to the first definition, then define

$$Z_i = \{a_i\} \cup \left(\bigcup_{t=1}^{\infty} X_i^t\right)$$

and $Z_j = \bigcup_{t=1}^{\infty} X_j^t$ for each $j \in N \setminus \{i\}$ (and define $X_i^1 = \varnothing$). This suffices to prove that the first definition implies the second!

If $a_i \in A_i$ is rationalizable according to the second definition, the define $\mu_i^1 = \mu_i(a_i)$ and $\mu_j^t = \mu_j(a_j)$ for $t \geq 2$ and $j \in N$. We now only need to define the X_j^ts. The set X_j^t for $t \geq 2$ is defined as the set of a'_j such that there is some player $k \in N \setminus \{j\}$, some action $a_k \in X_k^{t-1}$, and some a_{-k} in the support of $\mu_k(a_k)$ such that $a'_j = (a_{-k})_j$. The sets X_j^1 for $j \in N \setminus \{i\}$ are defined as the set of a'_j such that a_{-i} is in the support of $\mu_i(a_1)$ such that $a'_j = (a_{-i})_j$ (and $X_j^1 \neq \varnothing$ by convention). □

Remark. Every action used with positive probability by some player in a correlated equilibrium of a finite strategic game is rationalizable. However, the converse is not necessarily true – the set of rationalizable strategies is larger than the set of correlated equilibrium strategies.
Example. *Cournot Revisited* Consider the game with $N = \{1, 2\}$, $A_i = [0, 1]$, and

$$u_i(a_1, a_2) = a_i\left(1 - \sum_{j=1,2} a_j\right)$$

Player i's best response is $B_i(a_j) = \frac{1-a_j}{2}$, so the Nash equilibrium is $a_i = a_j = \frac{1}{3}$. Let's consider the set of rationalizable strategies $Z_i = Z_j = Z$ (by symmetry). Of course, $Z_i \subseteq A_i$. Define $m = \inf Z$ and $M = \sup Z$. A best response by i is a maximum of $a_i(1 - a_i - \mathbb{E}(a_j))$. Thus,

$$B_i(\mathbb{E}(a_j)) \in \left[\frac{1-M}{2}, \frac{1-m}{2}\right]$$

We need to have:

$$m \geq \frac{1-M}{2} \text{ and } M \leq \frac{1-m}{2} \Longrightarrow 2m \geq 1-M \geq 1-\frac{1-m}{2} = \frac{1+m}{2} \Longrightarrow m \geq \frac{1}{3}$$

and similarly, $M \leq \frac{1}{3}$, so $M = m = \frac{1}{3}$

Remark. In the definitions above, the beliefs of player i are presented as a general probability distribution over A_{-i}, meaning possible with correlated actions. An alternative is to assume that agents randomize in an independent way. These two are not equivalent.

Counterexample. Consider the game where player 3 selects the matrix M_i, where the number in each element is the common payoff of all players. Player 1 chooses the column, player 2 chooses the row, and they do not know which matrix they are in:

$$M_1 = \begin{bmatrix} 8 & 0 \\ 0 & 0 \end{bmatrix} \quad ; \quad M_2 = \begin{bmatrix} 4 & 0 \\ 4 & 0 \end{bmatrix} \quad ; \quad M_3 = \begin{bmatrix} 0 & 0 \\ 0 & 8 \end{bmatrix} \quad ; \quad M_4 = \begin{bmatrix} 3 & 3 \\ 3 & 3 \end{bmatrix}$$

In this game, M_2 is a rationalizable choice by player 3. We have that

$$U \in B_1(L, M_2) \quad ; \quad D \in B_1(R, M_2) \quad ; \quad L \in B_2(U, M_2) \quad ; \quad R \in B_2(D, M_2)$$

and

$$M_2 \in B_3\left(\frac{1}{2}(U, L) + \frac{1}{2}(D, R)\right)$$

so M_2 is rationalizable with $Z_1 = \{U, D\}$, $Z_2 = \{L, R\}$, and $Z_3 = \{M_2\}$. However, M_2 is not rationalizable if we require beliefs in which the actions of player 1 and player 2 are independent. Let p be the probability by which 1 selects U and q be the probability by which 2 chooses L. For M_2 to be optimal we need

$$4pq + 4(1-p)(1-q) \geq \max\{8pq, 8(1-p)(1-q), 3\}$$

which is impossible.

1.7 Dominance

We start with two related definitions:

Definition. An action of player i is a *never-best response* if it is not a best response to any belief of player i.

Remark. An action that is never a best response cannot be rationalized.

Definition. An action a_i of i is *strictly dominated* if there is a mixed strategy α_i of i such that

$$U_i(a_{-i}, \alpha_i) > U_i(a_{-i}, a_i)$$

for all a_{-i}.

Lemma 1.1. *An action of a player in a finite game is a never best response if and only if it is strictly dominated.*

Proof. Take the game $G = \langle N, \{A_i\}, \{u_i\}\rangle$ and let $a_i^\star \in A_i$. Consider the auxiliary strictly competitive game G' (see O&R Definition 21.1, it's a game where $a \succsim_1 b$ if and only if $b \succsim_2 a$) in which the set of actions of player 1 is $A_i \setminus \{a_i^\star\}$, the set for player 2 is A_{-i}, and the preferences of player 1 are represented by the payoff function $v_1(a_i, a_{-i}) = u_i(a_{-i}, a_i) - u_i(a_{-i}, a_i^\star)$. For any given mixed strategy profile (m_1, m_2) of G', we denote by $v_1(m_1, m_2)$ the expected payoff of player 1.

The action $a_i^\star$ is a never-best response in G if and only if for any mixed strategy of player 2 in G' there is an action of player 1 that yields a positive payoff (*i.e.* if and only if $\min_{m_2} \max_{a_i} v_1(a_i, m_2) > 0$). This is so if and only if $\min_{m_2} \max_{m_1} v_1(m_1, m_2) > 0$, by the linearity of v_1 in m_1.

By Nash's Theorem, the game G' has a mixed strategy Nash equilibrium, and by von Neumann's Minimax Theorem applied to the mixed extension of G', we have that $\min_{m_1} \max_{m_2} v_1(m_1, m_2) > 0$ if and only if $\max_{m_1} \min_{m_2} v_1(m_1, m_2) > 0$, which holds if and only if there exists a mixed strategy $m_1^\star$ of player i in G' for which $v_1(m_1^\star, m_2) > 0$ for all m_2 (meaning, for all beliefs on A_{-i}). Since $m_1^\star$ is a probability measure on $A_i \setminus \{a_i^\star\}$ it is a mixed strategy of player 1 in G; the condition $v_1(m_1^\star, m_2) > 0$ for all m_2 is equivalent to $U_i(a_{-i}, m_1^\star) - U_i(a_{-i}, a_i^\star) > 0$ for all $a_{-i} \in A_{-i}$, which is equivalent to $a_i^\star$ being strictly dominated. □

Example. *Iterated Deletion of Dominated Strategies* Consider the following game:

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

In this game, B is dominated by $\frac{1}{2}T + \frac{1}{2}M$. Then, without it, R strictly dominates L, and without L, M strictly dominates T. Thus, the game becomes $(3,1)$, attained with (M,R).

Let's formalize this:

Definition. The set $X \subseteq A$ of outcomes of a strategic game *survives iterated deletion of strictly dominated actions (idsds)* if $X = \bigtimes_{j\in N} X_j$ and there is a collection $\{\{X_t^j\}_{j\in N}\}_{t=0}^T$ of sets that satisfies the following conditions for each $j \in N$:

1. $X_j^0 = A_j$ and $X_j^T = X_j$; and $X_j^{t+1} \subseteq X_j^t \ \forall\, t \in \{0, \dots, T-1\}$
2. For each $t \in \{0, \dots, T-1\}$ every action of player j in $X_j^t \setminus X_j^{t+1}$ is *strictly dominated* in the game $\langle N, \{X_i^t\}, \{u_i^t\}\rangle$, where u_i^t is the function u_i restricted to $X^t = \bigtimes_{i\in N} X_i^t$.

Proposition 1.3. *If $X = \bigtimes_{j\in N} X_j$ survives iterated deletion of strictly dominated actions in a finite strategic game, then X_j is the set of player j's rationalizable actions for each j.*

Proof. $(Z_j \subseteq X_j)$: Assume a_i is rationalizable with supporting sets $\{Z_j\}_{j\in N}$. Then for any t we must have $Z_j \subseteq X_j^t$ since each action in Z_j is a best response in A_j to some belief over Z_{-j}, hence not strictly dominated in game $\langle N, \{X_j^t\}, \{u_j^t\}\rangle$.

$(X_j \subseteq Z_j)$: Every action in X_j is a best response to something in X_{-j}. However, we need to show that it's a best response to something in A_j. FSOC, assume that there is some $a_j \in X_j$ that is a best response in X_j^t to a belief μ_j in X_{-j} but not in X_j^{t-1}. Then there is some $b_j \in X_j^{t-1} \setminus X_j^t$ that is a best response in X_j^t to a belief μ_j in X_{-j}. However, by definition $b_j \in X_j^t$, which contradicts the assumption. □

Remark. This relies on the fact that an action is strictly dominated if and only if it is never a best response. It also relies on the fact that for rationalizability we may need beliefs to be correlated. If we require beliefs to be independent, X_j may not be rationalizable. This is only a problem with more than 2 players. For an example, look to the counterexample above – when we imposed independence, M_2 was no longer rationalizable, but it survives iterated deletion.

We can expand this concept further:

Definition. An action a_i is *weakly dominated* for player i if there is a mixed strategy α_i such that

$$U_i(a_{-i}, \alpha_i) \geq U_i(a_{-i}, a_i)$$

for all $a_{-i} \in A_{-i}$, with the inequality holding strictly for at least one element of A_{-i}.

An action that is weakly dominated (but not strictly dominated) is a (weak) best response to some belief. Eliminating these actions may not be rational, but there is no strict advantage to using them.

Definition. The set $X \subseteq A$ of outcomes of a strategic game *survives iterated deletion of weakly dominated strategies (idwds)* if $X = \bigtimes_{j\in N} X_j$ and there is a collection $\{\{X_t^j\}_{j\in N}\}_{t=0}^T$ of sets that satisfies the following conditions for each $j \in N$:

1. $X_j^0 = A_j$ and $X_j^T = X_j$ and $X_j^{t+1} \subseteq X_j^t$ for each $t \in \{0, \dots, T-1\}$
2. For each $t \in \{0, \dots, T-1\}$ every action of player j in $X_j^t \setminus X_j^{t+1}$ is *weakly dominated* in the game $\langle N, \{X_i^t\}, \{u_i^t\} \rangle$, where u_i^t is the function u_i restricted to $X^t = \times_{i \in N} X_i^t$.

Remark. The order of deletion now matters!

Example. Consider the game

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

If we eliminate T, then B, the outcome is (M, R). However, if we eliminate B and then R, (M, R) cannot be the equilibrium.

Definition. A strategic game is *dominance solvable* if all players are indifferent between all outcomes that survive the iterative procedure in which all the weakly dominated actions of each player are eliminated at each stage.

Example. Consider the game

	L	R
U	$1,0$	$0,0$
D	$0,1$	$0,0$

This game is dominance solvable, where we end at (U, L). However, if we eliminate D, then neither L nor R is dominated, so idwds gives us another outcome.

1.8 Supermodular Games

Remark. In a lot of applications, we have strategic complimentarities, meaning that one player's action is increasing in the other player's action. Games with this property have nice characteristics. First, recall some definitions from ECON 6170:

Definition. $u_i(s_i, s_{-i})$ has *increasing differences* in (s_i, s_{-i}) if for all $(s_i, \tilde{s}_i)$ and $(s_{-i}, \tilde{s}_{-i})$ such that $s_i \geq \tilde{s}_i$ and $s_{-i} \geq \tilde{s}_{-i}$ we have

$$u_i(s_i, s_{-i}) - u_i(\tilde{s}_i, s_{-i}) \geq u_i(s_i, \tilde{s}_{-i}) - u_i(\tilde{s}_i, \tilde{s}_{-i})$$

If all inequalities are strict, then we have strict increasing differences. Recall from ECON 6170 that if s_i, s_{-i} are elements of sublattices of $\mathbb{R}^n$, a function has increasing differences in (s_i, s_{-i}) if and only if it is *supermodular* in s_i for all fixed s_{-i}, meaning if

$$u_i(s_i, s_{-i}) + u_i(\tilde{s}_i, s_{-i}) \leq u_i(s_i \wedge \tilde{s}_i, s-i) + u_i(s_i \vee \tilde{s}_i, s-i)$$

Additionally, if $u_i \in \mathbf{C}^2$, then u_i is supermodular in $s_i \iff \frac{\partial^2 u_i}{\partial s_{i,j} \partial s_{i,k}} \geq 0$. We can also define supermodularity in a vector s as

$$u_i(s) + u_i(\tilde{s}) \leq u_i(s \wedge \tilde{s}) + u_i(s \vee \tilde{s})$$

for all $s, \tilde{s}$. Note that supermodularity in s implies increasing differences in (s_i, s_{-i}) and supermodularity in s_i.

Remark. With increasing differences, an increase in the strategies of the opponents raises the desirability of choosing a higher strategy.

Definition. A (resp. strictly) *supermodular game* is a game in which for each i, S_i is a sublattice[1] of $\mathbb{R}^{m_i}$, u_i has (resp. strictly) increasing differences in (s_i, s_{-i}), and u_i is (resp. strictly) supermodular in s_i (or

[1]Recall: A subset S of X is a sublattice if $x, y \in S \implies x \vee y \in S, x \wedge y \in S$.

equivalently: u_i is supermodular in s).

Example. *Bertrand Competition with Differentiated Products* Consider an oligopoly with linear demand functions

$$D(p_i, p_{-i}) = a_i - b_i p_i + \sum_{j \neq i} d_{ij} p_j$$

with $b_i > 0$ and $d_{ij} > 0$, and constant marginal cost c_i. This is a supermodular game, as profits are:

$$\pi_i(p_i, p_{-i}) = (p_i - c_i) D_i(p_i, p_{-i})$$

where, since $\pi_i \in \mathbf{C}^2$, we can easily verify that $\frac{\partial^2 \pi_i}{\partial p_i \partial p_j} \geq 0$, so the game has increasing differences (and trivially profit is supermodular since the strategy is unidimensional).

Example. *Diamond's Search Model* A player's utility depends on his search intensity and other players' intensities:

$$u_i(s_i, s_{-i}) = \alpha s_i \cdot \sum s_j - c(s_i)$$

Here, we have strategic complimentarities.

Example. *Solving Bertrand* Assume $N = 2$, $A_i = [0,1]$, and $D_i(p_i, p_j) = 1 - 2p_i + p_j$, meaning that we have

$$\pi_i(p_i, p_j) = p_i[1 - 2p_i + p_j]$$

and taking the first order condition we get that

$$\frac{\partial}{\partial p_i} \pi_i(p_i, p_j) = 1 - 4p_i + p_j = 0 \Longrightarrow p_i^\star = \frac{1 + p_j}{4}$$

Since $p_j \in [0,1]$, we can say that $p_i^\star \in [1/4, 1/2]$. More formally, we can say that any strategies outside that interval are dominated by either 1/4 or 1/2. We can again repeat this logic, however, since the game is symmetric! $X_1^i = [1/4, 1/2]$, so $X_2^i = [5/16, 3/8]$, and $X_3^i = [21/64, 11/32]$, and so on. We have that

$$\lim_{t \to \infty} X_t^i = \{1/3\}$$

So this game is solvable by iterated deletion of strictly dominated strategies.

In general, we have:

Theorem 1.5. *Let (S, u) be a supermodular game. Then:*

1. *The set of strategies surviving iterated strict dominance has greatest and least elements $\bar{a}, \underline{a}$*
2. *$\bar{a}$ and $\underline{a}$ are Nash equilibria*

Proof. Note first that defining the best response correspondence as $BR_i(a_{-i}) = \operatorname{argmax}_{a_i} u_i(a_i, a_{-i})$, by continuity and compactness we have that $BR_i(a_{-i})$ is nonempty and has a greatest and least element, $\overline{BR}_i(a_{-i})$ and $\underline{BR}_i(a_{-i})$ respectively. Furthermore, increasing differences implies that if $a'_{-i} > a_{-i}$, then $\overline{BR}_i(a'_{-i}) > \overline{BR}_i(a_{-i})$ and $\underline{BR}_i(a'_{-i}) > \underline{BR}_i(a_{-i})$.

Start from $A = A^0$ and let $\bar{a}^0 = (\bar{a}_1^0, \ldots, \bar{a}_n^0)$ be the largest element. Define $\bar{a}_i^1 = \overline{BR}_i(\bar{a}_{-i}^0)$. Then any $a_i > \bar{a}_i^1$ is strictly dominated by $\bar{a}_i^1$. We can iterate and obtain $\bar{a}_i^k$, and notice that the sequence is decreasing in k. Define $\bar{a}_i = \lim_{k \to \infty} \bar{a}_i^k$, and note that by continuity, $\bar{a}_i \in \overline{BR}_i(\bar{a}_{-i})$, so it is a Nash equilibrium. The same process works to find $\underline{a}$. □

We could have also used another fixed point theorem:

Theorem 1.6. Tarski *If S is a nonempty compact sublattice of $\mathbb{R}^m$ and $f : S \to S$ is non-decreasing, then f has a fixed point.*

Remark. When strategies are one-dimensional, Tarski can be applied almost immediately with increasing differences since $BR_i(a_{-i})$ is also a non-empty compact monotonic sublattice and it has a monotone selection.

When strategies are multi-dimensional, it remains to show that $f(S)$ is a sublattice.
Theorem 1.7. *Consider a supermodular game such that (i) S_i is a complete sublattice and bounded, and (ii) u_i is continuous and bounded. Then iterated deletion of strictly dominated strategies yields a set of strategies in which the greatest and least elements are Nash equilibria $\overline{s}$ and $\underline{s}$.*

Proof. Since S is a complete lattice, there is a greatest element $s^0 = (s_1^0, \ldots, s_I^0)$. Let s_i and s_i' be two strategies in $r_i^\star(s_{-i}^0)$ such that there is no $s_i'' \in r_i^\star(s_{-i}^0)$ such that $s_i'' > s_i$ or $s_i'' > s_i'$.

If $s_i \neq s_i'$, consider $s_i \wedge s_i'$. We have:

$$u_i(s_i, s_{-i}^0) - u_i(s_i \wedge s_i', s_{-i}^0) \leq u_i(s_i \vee s_i', s_{-i}^0) - u_i(s_i', s_{-i}^0) < 0$$

where the weak inequality follows from supermodularity and the strict inequality follows from $s_i \wedge s_i' > s_i'$, since by assumption there is no strictly larger element in the best response correspondence.

We have a contradiction, so we can assume that $r_i^\star(s_{-i}^0)$ has a single greatest element, which we call s^1. We can repeat this logic to create s^n.

Consider an element s_i where $s_i \not\leq s_i^n$. Then it is dominated by $s_i \wedge s_i^n < s_i$ (whenever $s_{-i} \leq s_{-i}^n$). To see this, note that

$$\begin{aligned} u_i(s_i, s_{-i}) - u_i(s_i \wedge s_i^n, s_{-i}) &\leq u_i(s_i, s_{-i}^{n-1}) - u_i(s_i \wedge s_i^n, s_{-i}^{n-1}) \\ &\leq u_i(s_i \vee s_i^n, s_{-i}^{n-1}) - u_i(s_i^n, s_{-i}^{n-1}) < 0 \end{aligned}$$

where the first inequality follows from increasing differences, the second from supermodularity, and the third from the fact that s_i^n is the greatest best response to s_{-i}^{n-1} and $s_i \vee s_i^n > s_i$. Then $\{s_i^n\}$ is bounded below and decreasing, so it converges to some $\bar{s}$.

It remains to show that $\bar{s}$ is a Nash equilibrium. Fix some s_i, and by optimality $u_i(s_i^{n+1} s_{-i}^n) \geq u_i(s_i, s_{-i}^n)$. Finally, by continuity,

$$u_i(\bar{s}_i, \bar{s}_{-i}) \geq u_i(s_i, \bar{s}_{-i})$$

We can similarly obtain $\underline{s}$ as the lower bound. ☐

2 Extensive Games

2.1 Extensive Games with Perfect Information

Remark. In many strategic situations we have (i) more information than $\langle N, \{A_i\}, \{u_i\}\rangle$, and (ii) this information may be useful to predict the type of strategic interaction. We will study how to describe these environments, and introduce solution concepts to exploit them. The key in this analysis is *order of play*
Definition. An *extensive game* is a detailed description of the sequential structure of a decision problem. Two cases are relevant:

1. *Perfect information*: Each player has knowledge of all previous events and decisions. Either players play sequentially, or there are decision nodes in which more than one player makes a decision.

2. *Imperfect information*: Players are not perfectly informed about previous events.

We will focus on the first, to begin with.
Definition. An *extensive game with perfect information* $\langle N, \mathcal{H}, P, \{u_i\}\rangle$ is the following: (i) a set N of players, (ii) A set $\mathcal{H}$ of *histories* (which are sequences) with the following properties: $\varnothing \in \mathcal{H}$, $\{a^k\}_{k=1}^K \in \mathcal{H}$ then $\{a^k\}_{k=1}^L \in \mathcal{H}$ for all $L \leq K$, and if an infinite sequence $\{a^k\}_{k=1}^\infty$ satisfies that $\{a^k\}_{k=1}^K \in \mathcal{H}$ for all K,

then $\{a^k\}_{k=1}^{\infty} \in \mathcal{H}$, (iii) a function $P : h \to N$ for $h \in \mathcal{H}$ that assigns to each non-terminal history a member of N (see below), and (iv) preferences over terminal histories $u_i : Z \to \mathbb{R}$.

A history is *terminal* if either it is infinite or if $\exists\, K$ such that $\{a^k\}_{k=1}^{K} \in \mathcal{H}$ but $\{a^k\}_{k=1}^{K+1} \notin \mathcal{H}$ for any a^{K+1}. An extensive game in which $\mathcal{H}$ is finite is called a *finite extensive game*.

Remark. The interpretation of this construct is as follows: (i) Each history corresponds to a node, (ii) after each history h player $P(h)$ chooses an action in the set $A(h) := \{a : (h, a) \in \mathcal{H}\}$, and (iii) the empty history is the initial history.

Definition. An extensive game can be represented by a *tree*, a connected graph with no cycles.Each node has exactly one predecessor, so a node is a complete description of all events that preceded it: not just a state, or complete physical situation.

Example. *Simple Extensive Game* Consider $\langle N, \mathcal{H}, P, \{u_i\}\rangle$, where $N = 2$, $\mathcal{H} = \{\varnothing, U, D, UL, UR\}$, $P(\varnothing) = 1$, $P(U) = P(D) = 2$, and $Z = \{UL, UR, D\}$. Finally, $u_1(UR) = 2$, $u_2(UR) = 1$, $u_1(UL) = u_2(UL) = 0$, and $u_1(D) = 1$, $u_2(D) = 2$. Visually, the tree is Figure 1.

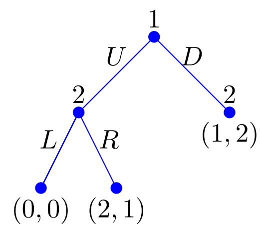

Figure 1: Simple Extensive Game

Definition. A *strategy* of a player i in an extensive game with perfect information is a function $s_i(h) \to A(h)$ for any $h \in H \setminus Z$ such that $P(h) = i$. A strategy specifies an action for any node in which a player is asked to choose an action. A *strategy profile* $s = (s_1, \ldots, s_N)$. For each strategy profile, an *outcome* $O(s)$ is the terminal node associated with the strategy profile. Note that we consider only pure strategies. If we consider mixes, then the outcome may be a distribution over terminal histories.

Definition. The *strategic form* of an extensive game with perfect information $\langle N, \mathcal{H}, P, \{u_i\}\rangle$ is the strategic game $\langle N, \{S_i\}, \{\tilde{u}_i\}\rangle$ in which S_i is the set of strategies in the extensive game, and $\tilde{u}_i(s) = u_i(O(s))$.

A *Nash equilibrium of an extensive game* is a Nash equilibrium of the associated strategic game.

Example. *Extensive Game with Strategic Form* Let's consider a modification to the previous example, in Figure 2.

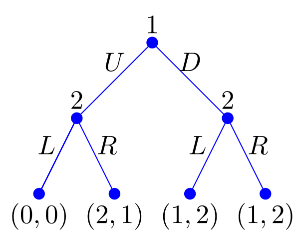

Figure 2: Modified Simple Extensive Game

The strategic form game is

	LL	LR	RL	RR
U	$(0, 0)$	$(0, 0)$	$(2, 1)$	$(2, 1)$
D	$(1, 2)$	$(1, 2)$	$(1, 2)$	$(1, 2)$

Here, LL stands for $s_2(U) = L$, $s_2(D) = L$; LR stands for $s_2(U) = L$, $s_2(D) = R$, and so on. Nash equilibria of the game are (U, RL), (U, RR), (D, LL), and (D, LR).

Definition. Define two strategies s_i and s_i' as *equivalent* if for each s_{-i} we have $u_i(s_i, s_{-i}) = u_i(s_i', s_{-i})$.

Definition. The *reduced form* of an extensive game is where we include only one member for each set of equivalent strategies. To wit:

	L	R
U	$(0,0)$	$(2,1)$
D	$(1,2)$	$(1,2)$

Remark. There are some problems with the Nash solution here. If player 1 chooses U, it's natural to think that player 2 will choose R. However, the equilibria (D, LL) and (D, LR) exist based on the conjecture that if player 1 chose U, player 2 would select L. That's clearly not going to happen, by rationality.

Definition. The *subgame* of an extensive game with perfect information Γ that follows from history h is the extensive form game $\Gamma(h) = \langle N, \mathcal{H} \mid_h, P \mid_h, \{u_i\} \mid_h \rangle$, where $\mathcal{H} \mid_h$ is the set of sequences h' of actions for which $(h, h') \in \mathcal{H}$, $P \mid_h$ is such that $P \mid_h (h') = P(h, h')$ for $(h, h') \in \mathcal{H}$, and $u_i(h'; h) \geq u_i(h''; h) \iff u_i(h, h') \geq u_i(h, h'')$.

Definition. A *subgame perfect equilibrium* is a strategy profile $s^\star$ in Γ in which for any history h the strategy profile $s^\star \mid_h$ is a Nash equilibrium of the subgame $\Gamma(h)$, where $s^\star \mid_h (h') = s^\star(h, h')$.

Example. *Stackelberg* Two firms 1 and 2 choose output levels $q_i \in [0, \infty)$. Firm 1 moves first, and the price is $p(q_1, q_2)$, so profit is $u_i(q_1, q_2) = q_i \cdot p(q_1, q_2) - c_i(q_i)$. To find a Nash equilibrium we find the reaction functions $r_2(q_1)$:

$$p(q_1 + r_2(q_1)) + p'(q_1 + r_2(q_1)) - c_1'(r_2(q_1)) = 0$$

and if we assume linear costs and demand we have that

$$r_i(q_{-i}) = \frac{1 - q_{-i} - c}{2}$$

Assume that 1 chooses first, and then 2. Now firm 1 optimizes knowing firm 2's reaction function, so they maximize

$$q_1 \cdot \left[1 - q_1 - \frac{1 - q_1 - c}{2}\right] - cq_1 = q_1 \left[\frac{1 - q_1 - c}{2}\right] - cq_1$$

From the FOC, $(1 - 2q_1 + c)/2 = c$, so $q_1 = (1 - c)/2$ and $q_2 = (1 - c)/4$.

Remark. To verify that a strategy $s^\star$ is a subgame perfect equilibrium, we need to check that for every $i \in N$ and every subgame $\Gamma(h)$, no strategy gives a strictly positive deviation. The following simplifies the calculation, by reducing the class of deviations we need to check.

Theorem 2.1. One-Shot Deviation Principle *In a finite extensive game with observed actions, a strategy profile s is a subgame perfect Nash equilibrium if and only if no player can strictly gain by deviating from s in a single stage and conforming to s thereafter.*

Proof. We want to show that s is a subgame perfect equilibrium if and only if there is no i and no $\hat{s}_i$ that agrees with s_i except at a single t and h^t and such that $\hat{s}_i$ is a better response to s_{-i} than s_i conditionally on h^t. The forward direction is immediate. We will focus on the backwards direction.

Proof by contrapositive: If s is not a subgame perfect equilibrium, then s violates the one-shot deviation principle. Suppose that s is not a subgame perfect equilibrium, meaning that there is a t and an h^t such that some i has a deviation $\hat{s}_i$ in the subgame $\Gamma(h^t)$. Let $\hat{t}$ be the largest t such that $\hat{s}_i(h^t) \neq s_i(h^t)$ (which exists because the game is finite). Consider an alternative strategy $\tilde{s}_i$ that agrees with $\hat{s}_i$ for all $t < \hat{t}$ and agrees with $s_i \mid_{h^t}$ from $\hat{t}$ on.

Since from any $h^{\hat{t}}$ it agrees with $s_i \mid_{h^{\hat{t}}}$ except for the first move, by the one-shot deviation principle this change can only increase the utility of i at any $h^{\hat{t}}$. Of course if the principle fails we are done, so assume that it does not. This means that $\tilde{s}_i$ is as good as $\hat{s}_t$ at h^t. If $\hat{t} = t + 1$, then $\tilde{s}_i = s_i$ and we have a contradiction. If $\hat{t} > t + 1$, then iterate the procedure until we have a contradiction. □

Remark. With no additional assumptions, this theorem fails in the infinite horizon case. Consider the following example, illustrated in Figure 3, where the payoff of playing infinite a is 1.

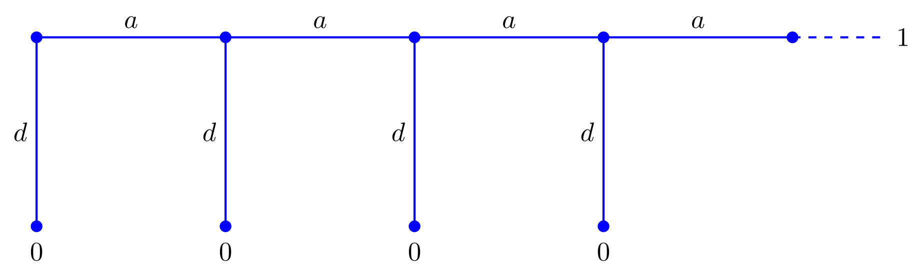

Figure 3: Infinite Game with no One-Shot Deviations

The strategy d after every history satisfies the one-shot deviation principle, but is clearly not a subgame perfect Nash equilibrium. To ensure that the one-shot deviation principle applies to infinite games, we need the following definition:

Definition. A game is *continuous at infinity* if for each player i the utility function $u_i(h)$ satisfies

$$\sup_{h,\hat{h} \text{ s.t. } h^t=\hat{h}^t} \left| u_i(h) - u_i(\hat{h}) \right| \to 0 \text{ as } t \to \infty$$

where $h, \hat{h}$ are infinite histories and $u_i(h), u_i(\hat{h})$ their respective utilities. Note that this condition is satisfied if the utilities are equal to a discounted sum of per-period payoffs $U_i^t(a^t)$ that are uniformly bounded.

Theorem 2.2. One-Shot Deviation Principle for Infinite Games *In an infinite horizon extensive game with observed actions that is continuous at infinity, a strategy profile s is a subgame perfect Nash equilibrium if and only if no player can strictly gain by deviating from s in a single stage and conforming to s thereafter.*

Theorem 2.3. Kuhn's Theorem *Every finite extensive game with perfect information has a subgame perfect equilibrium.*

Proof. (Constructive) We have a finite extensive game Γ with subgames $\{\Gamma(h)\}$, which are finite. Define $\ell(\Gamma(h))$ the length of the maximal history of $\Gamma(h)$. We now define $R(h)$ as a function which associates a terminal history h' to every history h such that $h' \succeq h$.

When $\ell(\Gamma(h)) = 0$, then we are in a terminal history and $R(h) = h$. Now assume we have defined $R(h)$ for all h with $\ell(\Gamma(h)) \leq k$. Consider an h' such that $\ell(\Gamma(h')) = k+1$. We have that $\ell(\Gamma(h', a)) \leq k$ for all $a \in A(h')$. Let $s_i(h')$ be such that $u_i(R(h', s_i(h'))) \geq u_i(R(h', a))$ for all $a \in A(h')$. Define $R(h) = R(h', s_i(h'))$. We have defined by induction $R(h)$ and a strategy s that is a subgame perfect equilibrium by the one-shot deviation principle. □

Remark. This method is called *backwards induction*. The intuition here is to solve the game from the end, from the most simple subgame.

Remark. We might also want to describe situations with some randomness – where nature also moves. This is easily incorporated.

Definition. An *extensive game with perfect information and chance moves* is a tuple $\langle N, \mathcal{H}, P, f_c, \{\succeq_i\}\rangle$, where now (i) P is a function from $\mathcal{H}$ to $N \cup \{c\}$ where c is for chance, (ii) for each h such that $P(h) = c$, $f_c(\cdot; h)$ is a probability distribution over $A(h)$, (iii) $\{\succeq_i\}$ are preferences over lotteries over terminal nodes.

Remark. We can also introduce a similar concept which introduces some uncertainty even in a game with perfect information and no chance.

Definition. An *extensive game with perfect information and simultaneous moves* is a tuple $\langle N, \mathcal{H}, P, \{u_i\}\rangle$ such that (i) N is the set of players, (ii) $\mathcal{H}$ is a sequence of $|P(h)|$ dimensional vectors of actions, (iii) P identifies the set of players who chose after history h, and (iv) u_i is the same as before.

A strategy is a function $s_i(h) \to A_i(h)$ for all $i \in P(h)$, and the definitions of subgames and subgame perfect equilibria apply here.
Remark. When we represent a game with simultaneous moves, we don't have perfect information. If we want to use a game tree representation, we need to describe this information. To this goal, we introduce:
Definition. *Information sets* are partitions of the histories with the interpretation that a player at a node x is unsure whether they are at x or any other $x' \in z(x)$. The same player must move at x and x', and we must have that $A(x) = A(x')$ for it to be true that $x, x' \in z(x)$.
Remark. Information sets can be used to describe information in a game tree. They could also describe situations in which information is degraded, meaning when a player might forget what they once knew. Games with perfect recall are games in which nobody forgets.
Remark. In a game with perfect information and simultaneous moves, we can generalize the one-shot deviation principle. However, we cannot guarantee the existence of a subgame perfect equilibrium in pure strategies. See: matching pennies.
Remark. A strategy specifies actions after nodes. Some histories can have zero probability given a player's strategy. We are requiring players to make choices even in situations that will never happen! We do this because it forms a basis for the beliefs of other players. The key assumption here is that rationality is still our guiding principle no matter what is observed.
Question. What happens when we end up in a history that has probability zero? What does that mean for the rationality of other players?

And then: What does this imply for how the other players *will* play? Perhaps they are irrational! That has implications for future play.

More generally: Past actions may be informative about how the opponents will play if there is some 'ambiguity' in the continuation subgame. Iterated deletion of weakly dominated strategies may capture this type of reasoning.
Example. *Battle of the Sexes (pt. 2)* Consider the following game, in extensive and normal form:

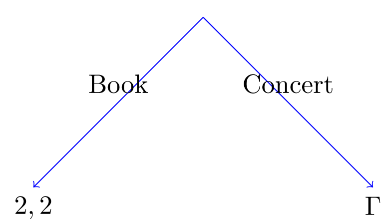

where:

$\Gamma \equiv$

	Bach	Stravinski
Bach	3,1	0,0
Stravinski	0,0	1,3

The subgames are (Book,S),S and (Concert,B),B. However, one is clearly more plausible than the other. We can see this in the strategic form of the full game:

	B	S
Book	$2, 2$	$2, 2$
B	$3, 1$	$0, 0$
S	$0, 0$	$3, 1$

Note that Book strictly dominates S for 1, and after that B weakly dominates S for 2. This implies (heuristically) that Bach is much more likely than Book.
Example. *BotS (Burning Money)* Consider the following game:

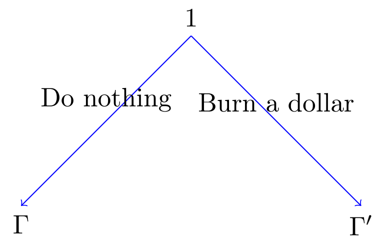

where

$\Gamma \equiv$

	Bach	Stravinski
Bach	3,1	0,0
Stravinski	0,0	1,3

and

$\Gamma' \equiv$

	Bach	Stravinski
Bach	2,1	-1,0
Stravinski	-1,0	0,3

where player 2 observes the choice to either do nothing or burn the dollar. We can solve this game with iterated deletion of strictly dominated strategies in the strategic form game:

	BB	BS	SB	SS
DnB	$3,1$	$3,1$	$0,0$	$0,0$
DnS	$0,0$	$0,0$	$1,3$	$1,3$
BdB	$2,1$	$-1,0$	$2,1$	$-1,0$
BdS	$-1,0$	$0,3$	$-1,0$	$0,3$

where DnB weakly dominates BdS, and after that is eliminated SB weakly dominates SS, so we end up with the unique pure-strategy Nash equilibrium in the game

	BB	BS	SB
DnB	$\boxed{3,1}$	$3,1$	$0,0$
DnS	$0,0$	$0,0$	$1,3$
BdB	$2,1$	$-1,0$	$2,1$

What is the logic here? In the original battle of the sexes, a player can guarantee a payoff of

$$\pi = \min_{\alpha \in [0,1]} \max\{3\alpha, 1-\alpha\} = \frac{3}{4}$$

In the original game, this is irrelevant. However, once the dollar has been burned, the only way for player 2 to guarantee this is by playing B. So if 1 burns the dollar they will play B, and if 1 does not burn the dollar 2 knows that they will play B, because otherwise they could do better by choosing to burn the dollar and guaranteeing 2.

Example. *The Centipede Game* Two players are in a process that they can alternatively stop or continue. At each time t, each player prefers stopping now to letting the opponent stop at $t+1$. In the last period $t = T-1$, the player prefers stop to continue. However, the terminal history T (attained by continuing at $T-1$) is better for both players than stopping at any $t < T-1$. One classical formulation is:

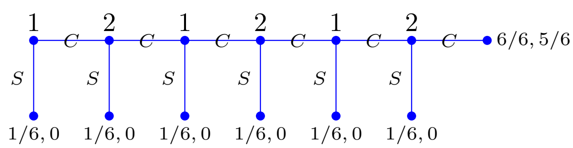

There is a unique subgame-perfect equilibrium: $s_i(h^t) = S$ for all i, t. Any pair of strategies in which player 1 chooses S in the first period and player 2 chooses S in the second is a Nash equilibrium. Is this realistic?
Remark. One way to reconcile the experimental observations is to note that cooperation is close to an equilibrium as long as the game is sufficiently long.
Definition. A profile $s^\star$ is an *ε-Nash equilibrium* if, for all players i and strategies s_i, we have

$$u_i(s^\star) \geq u_i(s_i, s^\star_{-i}) - \varepsilon$$

for some $\varepsilon > 0$.
Example. Consider the centipede game with T stages, so payoffs go to $T, T-1$, and normalize by dividing by T. Is cooperation up to k (for some k) optimal if T is sufficiently large? No deviation is optimal for $T \geq k$, since the strategy recommends to stop, and no deviation is optimal for $T \leq k-2$, since it is optimal to continue if the other player will continue. Finally, at $\tau = k-1$, the net benefit of a deviation is $1/T$, and for T sufficiently large $1/T < \varepsilon$. So this is an ε-Nash equilibrium.

2.2 Notable Dynamic Models

Model. *Rubinstein Bargaining* (from Rubinstein (1982)) Two players must split a pie of size 1. They alternate making offers: in even periods $t = 0, 2, 4, \ldots,$ Player 1 proposes $(x, 1-x)$ where x is the share allocated to Player 1. If Player 2 accepts, the payoffs are $(x, 1-x)$; and in odd periods $t = 1, 3, 5, \ldots,$ Player 2 proposes $(1-x, x)$, and Player 1 can accept or reject, and so on. We denote by x^i_j the amount allocated by player i to player j – so in Period 1, $x^1_1 = x$, and $x^1_2 = 1-x$. In each period, payoffs are discounted by (δ_1, δ_2) respectively, so payoffs for an allocation (x, y) in period t will be $(\delta_1^t x, \delta_2^t y)$.
Remark. There are many Nash equilibria in this game (in fact, every $x \in [0,1]$ admits a Nash equilibrium $(x, 1-x)$), but only one subgame perfect equilibrium. This restriction is what makes non-cooperative bargaining games tractable.

Consider the set of strategies: 1 always demands 1 and refuses anything less; 2 demands 0 and accepts anything. This is a Nash equilibrium (weakly, for 2, but still Nash), but is clearly not subgame perfect. In fact, any x proposed by 1 would be a Nash equilibrium with these same strategies. To show that it is not subgame perfect, observe: if 2 rejects the offer, they can offer $x \in (\delta, 1)$, where it is rational for 1 to accept since $u_1(x) = x > \delta = \delta u_1(1)$.

Here is a subgame perfect equilibrium (we will show that this is a subgame perfect equilibrium, and then that this is the unique subgame perfect equilibrium): Player i always demands a share

$$x^i_i = \frac{1-\delta_j}{1-\delta_i\delta_j}$$

when she makes an offer. Player i always demands

$$x^j_i = \frac{\delta_i(1-\delta_j)}{1-\delta_i\delta_j}$$

when she does not make an offer. First, we want to prove that this is a Nash equilibrium in any possible subgame. Usefully, there are only two subgames that are relevant – when i is the proposer and when i is the receiver. We will apply the one-stage deviation principle. All the other subgames are symmetric. Let's check the proposer deviations. Any proposal must satisfy

$$1 - \tilde{x}^i_i \geq x^i_j = \frac{\delta_j(1-\delta_i)}{1-\delta_j\delta_i} = 1 - \frac{1-\delta_j}{1-\delta_i\delta_j} = 1 - x^i_i \Longleftrightarrow \tilde{x}^i_i \leq x^i_i$$

Since i knows that x^i_i is accepted, it must be the case that $\tilde{x}^i_i \geq x^i_i$. Thus, this is either not profitable or not a deviation. Similarly, Player 2 refuses if

$$\tilde{x}_j^i < \delta_j x_j^j = \frac{\delta_j(1-\delta_i)}{1-\delta_i\delta_j} = x_j^i$$

and is willing to accept otherwise.

Rubinstein's key result is that this is the unique SPE. To see this, let $\bar{v}_i$ and $\underline{v}_i$ to be player i's supremum and infimum payoffs in the set of possible payoffs in a SPE. We must have that $\underline{v}_1 \geq 1 - \delta_2\bar{v}_2$, since 2 would always accept anything larger than $\delta_2\bar{v}_2$. Similarly, we must have that $\underline{v}_2 \geq 1 - \delta_1\bar{v}_1$. Moreover, we must have that

$$\bar{v}_1 \leq \max\{1 - \delta_2\underline{v}_2, \delta_1^2\bar{v}_1\}$$

The first inequality $\bar{v}_1 \leq 1 - \delta_2\underline{v}_2$ means that 2 rejects anything that gives her less than $\delta_2\underline{v}_2$, implying that $1 - x_1^1 \geq \delta_2\underline{v}_2$. The second follows from the fact that 1 can go for a rejected offer and wait one turn. So this implies that

$$\bar{v}_i \leq 1 - \delta_j\underline{v}_j$$

Combining the inequalities, we have that

$$\underline{v}_i \geq 1 - \delta_j\bar{v}_j \leq 1 - \delta_j(1-\delta_i)\underline{v}_i \Longrightarrow \underline{v}_i \leq \frac{1-\delta_j}{1-\delta_i\delta_j}$$

$$\bar{v}_i \leq 1 - \delta_j\underline{v}_j \geq 1 - \delta_j(1-\delta_i)\bar{v}_i \Longrightarrow \bar{v}_i \leq \frac{1-\delta_j}{1-\delta_i\delta_j}$$

So thus, we have that

$$\bar{v}_i = \underline{v}_i = \frac{1-\delta_j}{1-\delta_i\delta_j}$$

Remark. There are no mixed equilibria – even though the receiver is indifferent between accepting and rejecting, as soon as they decide to mix the best strategy for the sender will be to propose $\min\{(\delta_j v_j, 1]\}$, which is empty.

Remark. The unique payoff is determined by the discount factors and the order of play – the more patient player will attain higher payoff, and the first proposer will attain higher payoff in equilibrium. Note that the first mover advantage attenuates as $\delta \to 1$.

Definition. Often a game is played repeatedly over time. In this case, the game that is played repeatedly is called the *stage game* and the overall game is called the *repeated game*.

Remark. Even when this is done in finite horizons or the game has a unique equilibrium, this may lead to a larger set of equilibria. Repetitions allow the players to condition their actions on the actions taken by players in previous periods. In fact, even if the past actions are payoff irrelevant (meaning they do not affect the payoffs), conditioning on past actions makes the strategies *interactive* and thus more powerful. Equilibria may be associated to payoffs that are higher or lower for all players than the payoff in the unique equilibrium of the stage game.

Example. *Repeated Prisoner's Dilemma*. Recall:

	C	D
C	$(1,1)$	$(-1,2)$
D	$(2,-1)$	$(0,0)$

Defect is the unique equilibrium in the stage game. Note that past actions are payoff irrelevant – the way you played in the past does not affect the payoffs in the future. Consider the repeated version of the game, where strategies are functions of past actions a^t: $\sigma_i(a^t)$. If the game is repeated for T periods, we can write the payoff as

$$U_i = \frac{1-\delta}{1-\delta^{T+1}} \sum_{t=0}^{T} \delta^t u_i(\sigma(a^t))$$

where the first term allows us to express the payoffs as *average discounted payoffs*. As $T \to \infty$, we have that

$$U_i = (1-\delta)\sum_{t=0}^{\infty} \delta^t u_i(\sigma(a^t))$$

We claim that, as $T \to \infty$ when $\delta \geq 0.5$ there exists a subgame perfect equilibrium in which both players choose cooperate in equilibrium. An obvious SPE is playing defect forever. Consider the following *grim trigger* strategies. Player i plays C forever, but if player j plays D in some period t, Player i plays D for period $t+1$ and thenceforth forever.

To see that this is a subgame perfect equilibrium, consider deviations. If we are in the first state, playing C forever, then the payoffs for i are

$$u_i(C) = (1-\delta)\sum_{t=0}^{\infty} \delta^t = 1 \geq (1-\delta)[2+0+\cdots] = 2(1-\delta)$$

so as long as $\delta \geq 0.5$, C is weakly preferred. If we are in the second state, the strategy prescribes playing D forever, which is the stage game dominant strategy, so of course is an equilibrium. Note that we can obtain a strictly higher payoff in equilibrium than the stage game Nash strategies get.

Remark. The (possible) average payoffs of the two players can be represented in two dimensions:

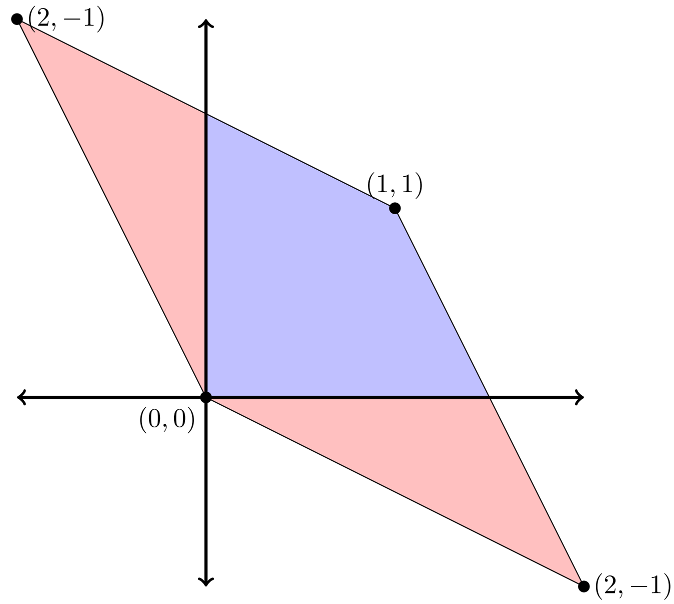

where the shaded areas show the convex hull of the stage payoffs, the blue shows the average payoffs attainable by a rationalizable strategy (the *Folk Theorem*).

Example. *Carrot and Stick* Consider the following game:

	A	B	C
A	2,2	2,1	0,0
B	1,2	1,1	-1, 0
C	0,0	0,-1	-1,-1

This game has the unique equilibrium (A, A), getting payoffs $(2, 2)$. However, consider the following SPE: In State I, Play B unless the other does something different. If they deviate, go to State II, where you play

C. If all others play C, return to State I. If not, stay in State II. Suppose we are in State I. Then payoffs for B are $(1-\delta)\sum_{t=0}^{\infty}\delta^t = 1$, while the best deviation yields $2-\delta+\delta^2+\delta^3+\cdots = 1+(1-\delta)(1-2\delta)$ so B is preferred as long as $\delta \geq 1/2$. In State II, the payoffs are:

$$u(C) = (1-\delta)[-1+\delta+\delta^2+\cdots] = 1-2(1-\delta)$$
$$u(B) < u(A) = (1-\delta)[0-\delta+\delta^2+\cdots] = 1+(1-\delta)(1-2\delta)$$

So C is optimal again as long as $\delta \geq 1/2$.
Example. *Multiple Equilibria in a Finite Game* Consider:

	A	B	C
A	0,0	3,4	6,0
B	4,3	0,0	0,0
C	0,6	0,0	5,5

This game has two pure equilibria: (B, A) and (A, B), and a mixed equilibrium $(3/7A+4/7B, 4/7A+3/7B)$. Note that all payoffs fall short of the maximal, $(5,5)$. In the twice-repeated game, we have a SPE that leads to the maximal payoff $(5,5)$: The strategies are to play C at $t=1$, and if (C,C) is attained play (B,A) at $t=2$. Otherwise, play the mixed equilibrium strategy. In equilibrium, payoffs are $(5,5)+\delta(4,3)$, and the gain from a deviation is (at most) 1 at $t=1$, with a loss of at minimum $\delta(3-12/7)$. The strategies are a SPE if $\delta(3-12/7) > 1 \iff \delta \geq 7/9$.
Example. *War of Attrition* Two animals are fighting for a prize with value v. The fighting cost is 1 per period. If an animal stops fighting at t, the opponent wins v, there is no fighting, and the game stops. There is a per-period discount factor δ. The payoff of quitting at time $\hat{t}$ is

$$L(\hat{t}) = -(1+\delta+\delta^2+\cdots+\delta^{\hat{t}-1})$$

the payoff of the winner is

$$W(\hat{t}) = -(1+\delta+\delta^2+\cdots+\delta^{\hat{t}-1})+\delta^{\hat{t}}v = L(\hat{t})+\delta^{\hat{t}}v$$

If both animals quit together, we assume the payoff is $L(\hat{t})$ for both.

As in the bargaining game, here we have several Nash equilibria. For example: i always fights, j always stops. This is a Nash equilibrium and is subgame perfect. Is the game indeterminate? If we look for a symmetric equilibrium, there is one unique one, in the form of stopping with probability p in each period where the game is continuing. It is easy to see that $p \in (0,1)$. In equilibrium, we must have that

$$L(t) = pW(t)+(1-p)L(t+1) \iff L(t)-L(t+1) = p[W(t)-L(t+1)]$$

where the left hand side is the payoff of stopping and the right is the payoff of continuing. Simplifying, we have that this becomes

$$\delta^t = p(\delta^t+\delta^t v) \iff p^\star = \frac{1}{1+v}$$

Remark. Note that both the symmetric and asymmetric equilibria we have seen so far are stationary. A stationary Nash equilibrium is always a Nash equilibrium, since subgames are all strategically equivalent.

3 Repeated Games

3.1 Folk Theorems

Let G be a normal form game with action spaces $A_1, \ldots, A_I$, payoff functions $g_i : A \to \mathbb{R}$, where $A = \bigtimes_i A_i$. Let $G^\infty(\delta)$ be the infinitely repeated version of G played at $t = 0, 1, 2, \ldots$ where players discount at δ and observe all previous actions. A history is $\mathcal{H}^t = \{a^0, a^1, \ldots, a^{t-1}\}$, and a pure strategy is $s_{i,t} : \mathcal{H}^t \to A_i$. The average discounted payoff is

$$u_i(a_i, s_{-i}) = (1-\delta) \sum_{t=0}^{\infty} \delta^t g_i(s_i(h^t), s_{-i}(h^t))$$

Our goal is to study the set of average payoffs that are associated to SPE of the repeated game as a function of δ. A few constraints immediately bound this set:
Definition. The *set of feasible payoffs* is the set of vectors $C \subseteq \mathbb{R}^I$

$$(v_1, \ldots, v_I) \in \text{Co}\{(v_1, \ldots, v_I) : \exists\, (a_1, \ldots, a_I) \text{ s.t. } g_i(a) = v_i \;\forall\, i\}$$

where Co$\{\cdot\}$ denotes the convex hull of $\{\cdot\}$. Naturally, the set of equilibria must be included in this set.

Another constraint is individual rationality:
Definition. A player's *min-max* payoff is

$$\underline{v}_i = \min_{s_{-i}} \max_{s_i} g_i(s_i, s_{-i})$$

where here s_i is a mixed strategy.
Definition. A payoff vector is *individually rational* if $v_i \geq \underline{v}_i \;\forall\, i$.
Lemma 3.1. *Any Nash equilibrium must be individually rational.*

Proof. Suppose that $s^\star$ is a Nash equilibrium, where for some i $v_i < \underline{v}_i$. Then we have that there exists some other s_i' that attains a strictly higher payoff for any s_{-i}, including $s_{-i}^\star$. Thus, $s_i^\star$ would not be a best response, and i would deviate for s_i'. □

We will start with the classic framework, which will highlight some of the key ideas. However, it's not nearly as nice as the SPE Folk result we will see later.
Theorem 3.1. Folk Theorem in Nash Equilibria *If $v = (v_1, \ldots, v_I)$ is feasible and strictly individually rational, then there exists $\delta^\star < 1$ such that for all $\delta > \delta^\star$, there is a Nash equilibrium of $G^\infty(\delta)$ with average payoffs $(v_1, \ldots, v_I)$.*

Proof. Assume that there exists a profile a such that $g_i(a) = v_i$ for all i. This is for simplicity and not without loss,[2] we will return to this later. Let m_{-j}^j be the strategy profile of players other than j that holds j to at most $\underline{v}_j$, and write m_j^j for j's best response to m_{-j}^j. Let $m^j = (m_j^j, m_{-j}^j)$. Now consider the following strategies:

- State I: Play a if there was no deviation or if there was more than one deviation
- State II: if j deviates, play m^j forever

We can verify this is a Nash equilibrium using one-stage deviation. If a is played, then j receives

$$(1-\delta)\left(v_j + \frac{\delta}{1-\delta} v_j\right) = v_j$$

[2]It does hold if we have either a continuum of actions or a randomization device.

If there is a deviation, then j receives

$$(1-\delta)\left(\bar{v}_j + \frac{\delta}{1-\delta}\underline{v}_j\right)$$

so deviation is not profitable if and only if

$$(1-\delta)(\bar{v}_j - v_j) \le \delta(v_j - \underline{v}_j)$$

As $\delta \to 1$, the left hand side goes to zero, so this condition holds for sufficiently large δ. Note that we are using the fact that v is *strictly* individually rational here. □

Remark. The issue here is that we are asking players to minimax after a deviation – but that might not be a Nash equilibrium of the subgame in question, so would not be a credible threat. See the example:

	A	B
A	8,8	0,-50
B	10,1	0,-50

Note here that $\underline{v}_1 = 0$ and $\underline{v}_2 = 1$. So despite the fact that $(8,8)$ is feasible and individually rational, the Nash Folk Theorem says we can achieve it as a Nash equilibrium, but the minimax threat is not credible (as 2 would get -50 forever!).

Theorem 3.2. SPE Folk Theorem (from Fudenberg and Maskin (1986)) *Let $V^\star$ be the set of feasible and strictly individually rational payoffs. Assume that* $\dim V^\star = I$. *Then for any $(v_1, \dots, v_I) \in V^\star$, there exists $\delta^\star < 1$ such that for any $\delta > \delta^\star$, there is a subgame perfect equilibrium of $G^\infty(\delta)$ with average payoffs $(v_1, \dots, v_I)$.*

Proof. Fixing a payoff vector $v \in V^\star$, we construct a SPE that achieves it. For convenience (and again with only a small loss), assume that there is a strategy profile a such that $g_i(a) = v_i$ for all i. Choose $v' \in \text{int}(V^\star)$ such that $\underline{v}_i < v'_i < v_i$ for all i. We choose N such that

$$\max_a g_i(a) + N\underline{v}_i < \min_a g_i(a) + Nv'_i$$

We choose $\varepsilon > 0$ such that for each i,

$$v'(i) = (v'_1 + \varepsilon, \dots, v'_{i-1} + \varepsilon, v'_i, v'_{i+1} + \varepsilon, \dots, v'_I + \varepsilon)$$

See the figure:

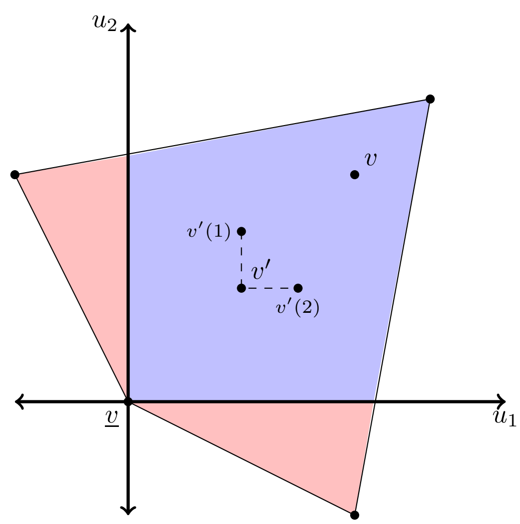

Assume further that there exists a^i such that $g(a^i) = v'(i)$. Assume that there is a pure strategy profile m^i that minimaxes i, so $g_i(m^i) = \underline{v}_i$. We will return to this assumption later. We now construct the following *carrot and stick* strategies:

- Stage I: Play a as long as nobody deviates. If j alone deviates, go to II_j (if two or more deviate, stay in I)
- Stage II_j: Play m^j for N periods, then go to III_j if nobody deviates. If k deviates, restart as II_k.
- Stage III_j: Play a^j as long as nobody deviates. If k alone deviates, go to II_k.

To check that these are optimal, we check every subgame in turn (for player i)

In subgame I, if i follows the strategy they get v_i, if they deviate they get

$$(1-\delta)(\max_a g_i(a) + \delta\underline{v}_i + \cdots + \delta^N\underline{v}_i + \delta^{N+1}v'_i + \cdots)$$

where deviation is obviously lower for sufficiently large δ since $\underline{v}_i < v'_i < v_i$.

In subgame II_i, suppose that there are $N' \leq N$ periods left. If I follows the strategy, they get

$$(1-\delta^{N'})g_i(m^i) + \delta^{N'}v'_i = q(N') = (1-\delta)g_i(m^i) + \delta q(N'-1)$$

where $g_i(m^i)$ is the payoff at the minimax strategy m^i for i. If i deviates, they do not improve in the deviating period and the punishment stage is restarted. So deviating attains

$$(1-\delta)g_i(m^i) + \delta q(N) < (1-\delta)g_i(m^i) + \delta q(N'-1)$$

In subgame II_j, suppose that there are $N' \leq N$ periods left. If i follows the strategy, they get

$$(1-\delta^{N'})g_i(m^j) + \delta^{N'}(v'_i + \varepsilon)$$

and if they deviate, they get

$$(1-\delta)\max_a g_i(a, m^j_{-i}) + \delta(1-\delta^N)\underline{v}_i + \delta^{N+1}v'_i$$

which is clearly strictly less.

Finally, in subgame III_i, if i follows the strategy, they get v'_i, and if they deviate they get

$$(1-\delta)\max_a g_i(a, a^i_{-i}) + \delta(1-\delta^N)\underline{v}_i + \delta^{N+1}v'_i$$

but this is strictly less since we assumed that N was such that $\max_a g_i(a) + N\underline{v}_i < \min_a g_i(a) + Nv'_i$. □

Remark. At two steps, we assumed that pure action profiles existed to generate the utilities we need. We could generate this if we had public randomizations, or by constructing strategies that change over time.

Remark. We also assumed (further) that the minimax strategy could be implemented with a pure strategy. If this is not the case, we need to ensure that players are willing to use a mixed minimax strategy. To this goal, a player i must be willing to mix over a set of actions. This is possible only if the player is indifferent among the actions. This is possible, by changing around the exact values of ε. However, it complicates the proof considerably.

Remark. The assumption here is that $\dim V^\star = I$. We could weaken this to no player having payoffs that are an affine transformation of another player's payoffs. A certain qualification on payoffs is, however, necessary. Consider the following game, where $P1$ selects rows, $P2$ selects columns, and $P3$ selects the matrix:

	A	B
A	$(1,1,1)$	$(0,0,0)$
B	$(0,0,0)$	$(0,0,0)$

A

	A	B
A	$(0,0,0)$	$(0,0,0)$
B	$(0,0,0)$	$(1,1,1)$

B

Visually, this is Figure 4.

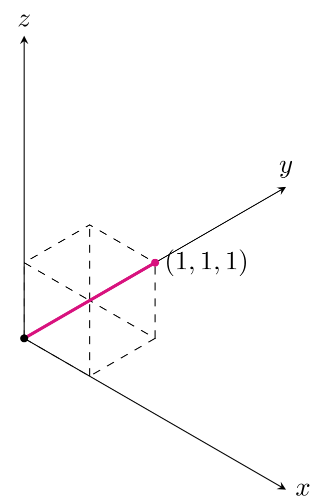

Figure 4: Folk Counterexample

In this game, the minimax is 0 for all players, and the set of feasible individually rational payoffs is $V^\star = \{(v,v,v) : v \in (0,1)\}$ (on the plot in magenta). Can we get all of these as SPE? No! Let

$$\underline{v} = \inf\{v : (v,v,v) \text{ is a SPE payoff}\}$$

For v to be a SPE we need that $v \geq \frac{1}{4}(1-\delta) + \delta\underline{v}$ since there must be at least two players in the three with $s_i(A) \geq 1/2$ or $s_i(B) \geq 1/2$ in the first period. Say that $s_1(A) \geq 1/2$ or $s_1(B) \geq 1/2$. Then $\underline{v} \geq \frac{1}{4}(1-\delta+\delta\underline{v} \iff \underline{v} \geq \frac{1}{4}$ since 3 can choose A in the first period. Therefore, there is no SPE with payoffs, say, $(1/8, 1/8, 1/8)$.

3.2 Imperfect Public Monitoring

A limitation of the repeated games model studied so far is that we assume that actions are observable. In many interesting applications, this is not the case. Imagine a case where actions are unobservable, but there are *imperfect public signals* correlated to the actions. These signals can be used in a repeated game to sustain cooperation (or more generally as inputs to a player's strategies).

Model. *Imperfect Public Signals* Let $(A_1, \ldots, A_I)$ be finite actions sets, and let Y be a finite set of public outcomes. Let $\pi(y \mid a) = \mathbb{P}(y \mid a)$. Let $r_i(a_i, y)$ be i's payoff if she plays a_i and the public outcome is y. Player i's expected payoff is

$$g_i(a) = \sum_{y \in Y} \pi(y \mid a) \cdot r_i(a_i, y)$$

A mixed strategy is $\alpha_i \in \Delta(A_i)$. Payoffs are defined the obvious way. The *public information* at the start of period t is $h^t = (y^0, \ldots, y^{t-1})$, and the *private information* for player i in period t h_i^t is her sequence of past actions. A strategy for i is a sequence of maps $\sigma_i^t : (h^t, h_i^t) \to \Delta(A_i)$.

Definition. A *public strategy* for player i is a sequence of maps $\sigma_i^t : h^t \to \Delta(A_i)$.

We focus on public strategies because they are simple and lead to a nice structure for the game.[3] Player i's average discounted payoff for the game if she gets a sequence of payoffs $\{g_i^t\}$ is

$$(1-\delta)\mathbb{E}_\sigma \sum_{t=0}^{\infty} \delta^t g_i(\sigma(h^t))$$

Definition. A profile $(\sigma_1, \dots, \sigma_I)$ is a *perfect public equilibrium (PPE)* if: (i) σ_i is a public strategy for all i, and (ii) for each date t and *public* history h^t, the strategy is a Nash equilibrium starting from that point.
Remark. A player might be uncertain to which node they are at – since we have *imperfect* public information, officially this is distinct from SPE (and SPE has no bite here). However, since opponents don't use private information in their own strategies, all possible nodes have the same distribution over opponent play, so there's no need to distinguish. Like the set of subgame perfect equilibrium payoffs in a repeated game model with perfect monitoring, the set of PPE payoffs is stationary.

A special case of public monitoring is when $Y = A$, and $\pi(y \mid a) = 1$. In that case, all PPE are SPE and vie versa. Note that PPE are perfect Bayesian equilibria of the repeated game, but not all perfect Bayesian equilibria are PPE.
Example. *Canonical Example* (from Green & Porter, 1984) Actions are interpreted as quantities, $a_i = q_i \in [0, Q]$, where quantities are unobserved. They determine an observed marker price $p = P(q, \varepsilon)$, where p is a random variable: $\lambda(q) = \mathbb{P}\{p \geq \hat{p} \mid q\}$. Green and Porter study collusion, an equilibrium in *trigger price strategies*. In Phase I, we produce $\hat{q}$. If $p \geq \hat{p}$, stay in this phase. If not, go to Phase II, where we play a static equilibrium for T periods. The value for these strategies in Phase I is:

$$\hat{v} = (1-\delta)g(\hat{q}) + \delta\left[\lambda(\hat{q}) + (1-\lambda(\hat{q}))\delta^T\right]\hat{v} \iff \hat{v} = \frac{(1-\delta)g(\hat{q})}{1-\delta\left[\lambda(\hat{q}) + (1-\lambda(\hat{q}))\delta^T\right]}$$

where we normalize the payoff of the punishment phase to zero. Obviously in Phase II, the static equilibrium is incentive compatible. In Phase I, we need that

$$(1-\delta)g(q_i, \hat{q}_{-i}) + \delta\left[\lambda(q_i, \hat{q}_{-i}) + (1-\lambda(q_i, \hat{q}_{-i}))\delta^T\right]\hat{v} \leq \hat{v}$$

It can be shown that there exist parameters under which these equilibria are attainable. A hypothetical designer designing a collusive equilibrium would choose $\hat{q}, \hat{p}, T$ to maximize $\hat{v}$ subject to the strategies being an equilibrium.
Model. *Dynamic Programming* (from Abreu, Pearce, & Stochetti, 1986 and 1990)
Definition. A pair (α, v) is *enforceable* with respect to δ and $W \subseteq \mathbb{R}^I$ if there exists a function $w : Y \to W$ such that for all i,

$$v_i = (1-\delta)g_i(\alpha) + \delta \sum_{y \in Y} \pi(y \mid \alpha) \cdot w_i(y)$$

and

$$\alpha_i \in \underset{\alpha_i' \in \Delta(A_i)}{\operatorname{argmax}} \left[(1-\delta)g_i(\alpha_i', \alpha_{-i}) + \delta \sum_{y \in Y} \pi(y \mid \alpha_i', \alpha_{-i}) \cdot w_i(y)\right]$$

Remark. The first condition says that the target payoff v can be decomposed into today's payoff and the expected continuation payoff, and that the strategy maximizes that decomposed value function. The second condition is incentive compatibility.

These conditions are similar to Bellman's Equation.
Definition. Let $B(\delta, W)$ be the set of payoffs v such that for some α, (α, v) is enforced with respect to δ

[3] This is just a refinement of the equilibrium concept – an equilibrium in public strategies is still a Nash equilibrium.

and W. Then $B(\delta, W)$ is the payoff set *generated* by δ, W.
Definition. Let $E(\delta)$ be the set of *PPE payoffs*
Proposition 3.1. $E(\delta) = B(\delta, E(\delta))$

Proof. ($\supseteq$): Fix $v \in B(\delta, E(\delta))$. Pick $w : Y \to E(\delta)$ such that w enforces (α, v). Now consider the following strategies: In period 0, play α. Then starting in period 1, play the perfect public equilibrium that gives payoffs $w(y_0)$. This is a PPE, so $v \in E(\delta)$.

($\subseteq$): If $v \in E(\delta)$, then there exists a PPE that gives v as payoffs. Suppose in this PPE, play in period 0 is α, and continuation payoffs are $w(y_0) \in E(\delta)$, since continuation corresponds to PPE play. The fact that nobody wants to deviate means that (α, v) is enforced by w, so $v \in B(\delta, E(\delta))$. □

Abreu, Pearce, and Stacchetti call this *factorization*. If it is possible to sustain average payoffs in W by promising different continuation payoffs in W, then W is self-generating. Formally,
Definition. W is *self-generating* if $W \subseteq B(\delta, W)$.
Remark. Note that $E(\delta)$ is self-generating. The set of static Nash equilibrium payoffs is also self-generating.
Proposition 3.2. *If W is self-generating, then $W \in E(\delta)$.*

Proof. Fix $v \in W$. Then $v \in B(\delta, W)$ so there is some $w : Y \to W$ and some α such that (α, v) is enforced by w. We will construct an equilibrium that gives v. In period 0, play α, and for an outcome y_0 set $v_1 = w(y_0)$. Then $v_1 \in W \subseteq B(\delta, W)$, so again there is some α_1 and some $w_1 : Y \to W$ such that (α_1, v_1) is enforced by w_1. Continue this strategy forever, to obtain the recommended strategies. After each public history there are no profitable deviations, and by construction the payoff is v. □
Corollary 3.1. *$E(\delta)$ is the largest self-generating set.*
Example. *Prisoner's Dilemma with Perfect Monitoring* We noted that games with perfect monitoring are special examples of games with public information, so we can think of the strategies above here. Let $Y = \{(C, C), (C, D), (D, C), (D, D)\}$. We will show that for $\delta \geq \frac{1}{2}$, the set $W = \{(0, 0), (1, 1)\}$ is self-generating. To this, we show that $(0, 0), (1, 1) \in B(\delta, W)$ for $\delta \geq \frac{1}{2}$. Seeing that $(0, 0) \in B(\delta, W)$ is trivial, since it is the static Nash outcome. In fact, for any δ and a_i we have that

$$0 \geq (1 - \delta) g_i(a_i, D) + \delta w_i(a_i, D)$$

Now consider $(1, 1)$. We will show that the strategy profile (C, C) and payoff profile $(1, 1)$ are enforced by $\delta \geq \frac{1}{2}$ and W. Let $w(C, C) = (1, 1)$ and $w(y) = (0, 0)$ for all $y \neq (C, C)$. Then

$$1 = (1 - \delta) g_i(C, C) + \delta w_i(C, C)$$

and for any a_i and $\delta \geq \frac{1}{2}$,

$$1 \geq (1 - \delta) g_i(a_i, C) + \delta w_i(a_i, C)$$

So $W \subseteq B(\delta, W)$ for $\delta \geq \frac{1}{2}$, meaning that W is self-generating.
Example. *Another PD Example* Consider this game:

	C	D
C	$(2, 2)$	$(-1, 3)$
D	$(3, -1)$	$(0, 0)$

Again assume that $Y = \{(C, C), (C, D), (D, C), (D, D)\}$. We now intend to prove that if $\delta \geq 1/3$, then $W = \{v, \hat{v}\}$ is self-generations, where

$$v = \left(\frac{3 - \delta}{1 + \delta}, \frac{3\delta - 1}{1 + \delta} \right) \qquad \text{and} \qquad \hat{v} = \left(\frac{3\delta - 1}{1 + \delta}, \frac{3 - \delta}{1 + \delta} \right)$$

Since it is a symmetric game and the continuations in W are permutations, we need only to show that v can be enforced with continuation in W.

Let the action profile α corresponding to v be (D, C) and the continuation payoffs be $w(D, C) = w(C, C) = \hat{v}$ and $w(D, D) = w(C, D) = v$. If players follow α, then the payoffs are

$$(1-\delta) \cdot \begin{pmatrix} 3 & -1 \end{pmatrix} + \delta \cdot \hat{v} = \left(\frac{3(1-\delta^2+3\delta^2-\delta)}{1+\delta}, -\frac{(1-\delta^2)+3\delta-\delta^2}{1+\delta} \right) = v$$

Clearly D maximizes the first player's action, since the current action does not affect future payoffs. If player 2 plays C as required by α, then they attain payoff $v_2 = \frac{3\delta-1}{1+\delta}$. If player 2 plays D, the payoff is 0 today and v_2 tomorrow. Thus, playing C is optimal as long as $v_2 \geq 0 \iff \delta \geq \frac{1}{3}$.

Remark. Recall that Folk Theorems aim to prove that all feasible and individually rational payoffs (*i.e.* $V^\star$) are achievable in equilibrium. A reasonable approximation is that any closed subset of $V^\star$ is achievable in equilibrium. Can we prove this statement with imperfect public monitoring?

If nothing is observed, then only the static Nash payoffs are attainable. It is reasonable to assume that the signal structure is sufficiently rich to provide incentives using expected payoffs.

Define $\pi(a_i \mid \alpha_{-i})$ to be a vector of probabilities on Y generated by a_i given α_{-i}, so it is a $|Y|$-dimensional vector. Define $\Pi(\alpha_{-i})$ to be the $|A_i| \times |Y|$-dimensional matrix that stacks the $\pi(a_i \mid \alpha_{-i})$. If we ignore feasibility constraints, we can implement a utility vector k as long as we can solve the system

$$(1-\delta)G(\alpha_{-i}) + \delta\Pi_i(\alpha_{-i})w_i = k$$

where $G_i(\alpha_{-i})$ is a column vector with generic element $g_i(a_i \mid \alpha_{-i})$ for all $a_i \in A_i$. The above system is solvable in w_i if $\Pi_i(\alpha_{-i})$ is full-rank.

Lemma 3.2. *The* individual full rank condition *is satisfied by a profile α if for each player i, $\Pi_i(\alpha_{-i})$ is invertible, meaning that the vectors $\pi_i(a_i \mid \alpha_{-i})$ are linearly independent.*

Remark. The individual full rank condition is not sufficient for a Folk Theorem. Consider the following example, from Radner, Myerson, & Maskin (1986). Two players can either work or shirk, at costs of 1 and 0 respectively. Output can be high or low, with probabilities from each outcome. If output is high, both players receive 4, otherwise they receive 0. The probability of high output is:

$$\pi_H(W, W) = \frac{9}{16} \qquad ; \qquad \pi_H(W, S) = \pi_H(S, W) = \frac{3}{8} \qquad ; \qquad \pi_H(S, S) = \frac{1}{4}$$

The individual full rank condition is satisfied at $\alpha = (W, W)$, since

$$\Pi_i(\alpha_{-i}) = \begin{bmatrix} \pi_H(W,W) & 1-\pi_H(W,W) \\ \pi_H(S,W) & 1-\pi_H(S,W) \end{bmatrix} = \underbrace{\begin{bmatrix} 9/16 & 7/16 \\ 3/8 & 5/8 \end{bmatrix}}_{\text{Full Rank}}$$

Despite the fact that the individual full rank condition holds, the Folk Theorem does not. To see this, let $v^\star$ be the highest payoff in any symmetric equilibrium. If the Folk Theorem is true, we should be able to get a payoff close to $4 \cdot \frac{9}{16} - 1 = \frac{5}{4} > 1$, since if the players choose (H, H), expected payoff is $(5/4, 5/4)$. We show that these payoffs cannot be approximated (in pure strategies only. The proof for mixed strategies follows fairly easily).

If the theorem holds, $v^\star > 1$, and since the equilibrium is stationary in equilibrium the players must choose (H, H) with at least positive probability. So we have that

$$v^\star = (1-\delta)\left[\frac{9}{16} \cdot (4 + \delta \cdot v_g) + \frac{7}{16} \cdot (0 + \delta \cdot v_b) - 1\right] \geq (1-\delta)\left[\frac{3}{8} \cdot (4 + \delta \cdot v_g) + \frac{5}{8} \cdot (0 + \delta \cdot v_b) - 1\right]$$

which implies that $v_g - v_b \geq \frac{4}{3} \cdot \frac{1-\delta}{\delta}$. However, by definition $v^\star \geq v_g$, so

$$v^\star \leq (1-\delta)\frac{5}{4} + \delta\left[\frac{9}{16} \cdot v^\star + \frac{7}{16} \cdot \left(v^\star - \frac{4}{3} \cdot \frac{1-\delta}{\delta}\right)\right] \iff v^\star \leq \frac{3}{2} \leq 1 \Rightarrow\Leftarrow$$

From this example, we learn that we need an additional assumption.

Define the matrix $\Pi_{i,j}(\alpha)$ to be the matrix formed by vertically concatenating the matrices $\Pi_i(\alpha_{-i})$ and $\Pi_j(\alpha_{-j})$. It is a $(|A_i| + |A_j|) \times |Y|$-matrix.

Lemma 3.3. *The* pairwise full-rank condition *is satisfied at action α for players i and j if $\Pi_{i,j}(\alpha)$ has maximal rank (equivalent to full column rank).*

Note that $\Pi_{i,j}(\alpha)$ cannot have full row rank, so the $|A_i| + |A_j|$ vectors admit at least one linear dependency. To see this, note that

$$\pi(\alpha) = \sum_{a_i \in A_i} \alpha_i(a_i) \cdot \pi(a_i \mid \alpha_{-i}) = \sum_{a_j \in A_j} \alpha_j(a_j) \cdot \pi(a_j \mid \alpha_{-j})$$

So we have that

$$\pi(a_1 \mid \alpha_{-i}) = \sum_{a_j \in A_j} \frac{\alpha_j(a_j)}{\alpha_1(a_1)} \cdot \pi(a_j \mid \alpha_{-j}) - \sum_{a_i \in A_i} \frac{\alpha_i(a_i)}{\alpha_1(a_1)} \cdot \pi(a_i \mid \alpha_{-i})$$

which is the linear dependency. Thus, we have:

Proposition 3.3. Imperfect Public Monitoring Folk Theorem *Suppose that $\dim V = I$, and both the individual full-rank condition (Lemma 3.2) and the pairwise full-rank condition (Lemma 3.3) hold. Then for any closed set $W \subset \int(V^\star)$, there exists some $\delta^\star < 1$ such that for any $\delta \geq \delta^\star$, $W \subset E(\delta)$.*

Remark. Some limitations exist here. A necessary condition to satisfy Lemma 3.3 is that $|A_i|+|A_j|-1 \leq |Y|$, which may be demanding. Indeed, it is not satisfied in the earlier Radner, Myerson, & Maskin example, in which we have two signals but $|A_i| + |A_j| - 1 = 3$. Moreover, the signal structure needs to be rich enough. For example, even if we have more than two signals it fails at symmetric profiles (*i.e.* (W, W)):

$$\Pi_{ij}(\alpha) = \begin{bmatrix} \pi_H(W,W) \\ \pi_H(S,W) \\ \pi_H(W,W) \\ \pi_H(W,S) \end{bmatrix}$$

where $\pi_H(W,W)$ has $|Y|$ dimensions. This matrix has rank $2 < 3$, since $\pi_H(L,H) = \pi_H(H,L)$.

3.3 Imperfect Private Monitoring

Example. Consider the following simply two-player game. In the first period, the players play a standard Prisoner's Dilemma (with payoffs $(1,1), (2,-1), (-1,2), (0,0)$). In the second period, they play the coordination game

	G	B
G	(k,k)	$(0,0)$
B	$(0,0)$	$(1,1)$

with $k > 2$. We can use the multiplicity of equilibria in the second game to incentivize cooperation in the first period. In perfect monitoring, the obvious strategy to cooperate, and play G if both cooperate, sustains cooperation for sufficiently large δ.

Now suppose that the first-period actions (a_1, a_2) are not observed. Rather, each player i observes a signal

$y_i \in \{c, d\}$ about her opponent's action. Suppose that

$$\mathbb{P}\{y_i = c \mid a_j\} = \begin{cases} 1 - \varepsilon & a_j = C \\ \varepsilon & a_j = D \end{cases}$$

where if ε is small, monitoring is almost perfect. We would expect that when ε is arbitrarily small, we could obtain cooperation. However, even for extremely small but positive ε, no pure strategy equilibrium where the players cooperate in period 1 can be sustained.

Observe that in the second period, i will want to play G if and only if she assigns probability $\frac{1}{k+1}$ or greater to the other player playing G. Consider strategies that call for each player to play C in the first period and G in the second if and only if $y_i = c$. If i plays C in the first period, she assigns probability $1 - \varepsilon$ to j observing c and hence to j playing G. However, regardless of what signal she observes she will want to play G, so won't want to follow the strategy.

Remark. Note that i and j's are *conditionally independent*. So long as i cooperates in the first period, she

GS: finish, with slide 14