

1. Rationalizability and Dominance

Defining a Game

A **normal form (strategic) game** $\langle N, (A_i), (u_i) \rangle$ consists of:

- A finite set of players N
- For each player $i \in N$, a non-empty set of actions A_i
- For each player $i \in N$, a utility function $u_i : \prod_{i=1}^N A_i \rightarrow \mathbb{R}$ ($u_i : A \rightarrow \mathbb{R}$)
 - Represents preferences over outcomes; we assume $u_i(\cdot)$ captures all concerns
 - When there is uncertainty, we assume that players' preferences can be represented by an expected utility function ($u_i(\cdot)$ represents preferences over the set of lotteries on $\prod_{i=1}^N A_i$ and is a linear combination of the probability of an outcome times the utility of that outcome for sure)

A **mixed strategy** for player i is any $\sigma_i \in \Delta(A_i)$ (the set of probability distributions over A_i)

- The set of mixed strategies includes pure strategies
- A **completely mixed strategy** is one where every $a_i \in A_i$ has some positive probability
- We assume that the players' mixed strategies are independent randomizations

A **belief** of player i , μ_i , is a probability measure on A_{-i} (the set of actions by the other players)

- A belief can be anything, including mixed strategies or correlated strategies between other players

An action $a_i \in A_i$ is a **best response** to μ_i if there is no other action that yields a strictly higher payoff to player i given his belief

- There can be more than one best response to a given μ_i

Rationalizability

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 (i) $a_i \in Z_i$ and (ii) every action $a_j \in Z_j$ is a best response to a belief μ_j of player j whose support is a subset of Z_{-j}

- There can only be positive belief probability on action that are a best response for the other player (definition is self-referencing)
- Given the set of actions that are rationalizable for you, what is rationalizable for me?

Dominance

An action $a_i \in A_i$ is **strictly dominated** in the strategic game $\langle N, (A_i), (u_i) \rangle$ if there exists a mixed strategy σ_i such that $u_i(\sigma_i, a_{-i}) > u_i(a_i, a_{-i})$ for all $a_{-i} \in A_{-i}$

- Some σ_i gives strictly better payoffs than a_i for every combination of other player actions

Iterated Elimination of Strictly Dominated Strategies (IESDS)

- In stage 1, eliminate a strictly dominated pure strategy for a player. In stage 2, consider only the remaining pure strategies and eliminate again. Continue until there are no more remaining strictly dominated pure strategies. The order of elimination does not matter.

Action a_i survives IESDS $\Leftrightarrow a_i$ is rationalizable

- For $N > 2$ games, requires that beliefs can be anything, as defined above (including beliefs about correlated strategies between other players); if we restrict beliefs to exclude correlated beliefs, then the set of rationalizable actions is smaller and the relationship with IESDS doesn't hold

An action $a_i \in A_i$ is **weakly dominated** in the strategic game $\langle N, (A_i), (u_i) \rangle$ if there exists a mixed strategy σ_i such that $u_i(\sigma_i, a_{-i}) \geq u_i(a_i, a_{-i})$ for all $a_{-i} \in A_{-i}$ and $u_i(\sigma_i, a_{-i}) > u_i(a_i, a_{-i})$ for some $a_{-i} \in A_{-i}$

- Some σ_i gives at least the same payoffs than a_i , with better payoff in at least one circumstance

A strategic game is **dominance solvable** if all players are indifferent between outcomes of the pure strategies that survive IEWDS

- For IEWDS, all weakly dominated pure strategies must be removed in a given stage (order matters)
- Never use IEWDS for anything other than dominance solvability

Finding Rationalizable Actions

- * Always use IESDS (remove strictly dominated actions of all players in any order)
- * Actions may be strictly dominated by mixed strategies
- * If an action is never a best response to another pure strategy, that can be a clue that it is strictly dominated (e.g., by a mixed strategy)
- * Do not try to remove non-rationalizable actions (This would require considering an uncountably infinite number of beliefs. If a pure strategy is strictly dominated, it is worse for any other player action so it's also worse for any mixture.)

Showing a Game is Dominance Solvable

- ✱ Remove all weakly dominated actions of all players simultaneously before moving to the next stage (order matters!) – if all players are indifferent at the end, the game is dominance solvable

Extensive Form Games

- Are drawn sequentially (as game trees)
- Map terminal histories to payoffs
- Include more information than normal form games (a normal form game can correspond to multiple extensive form games)
- Include nodes associated with all player actions and with all game histories (including the null history)
- Include information sets (the set of nodes that the acting player can tell the difference between)
 - If a player can't distinguish between two or more of his nodes, they are connected with a dotted line or a circle around the nodes
 - Within an information set, the actions that the player can take must be identical (otherwise he could tell the difference)
- Can include nodes assigned to “chaos” with a probability measure over the set of possible outcomes

Converting Extensive Form to Normal Form

- ✱ For each player, make a list of their information sets, including the associated history and their potential actions (list must include all actions, even ones with only one possible choice)
- ✱ Examples: $\emptyset : \{A, B\}$ $RrR : \{r, d\}$
 $C/D : \{C, D\}$ (where histories C and D are part of the same information set)
- ✱ Create a list of pure strategies for each player – each strategy will include an action for each listed information set (Example: $S_1 = \{RR, RD, DR, DD\}$)
- ✱ The order of the list of information sets will determine the order of the elements of the strategies (move left to right then top to bottom through the game for consistency)
- ✱ Create a normal form game with all pure strategies and their associated payoffs

Actions \neq Strategies

Actions are observable (the action that a player actually took)

Normal Form Games:

- Pure strategies = actions
- u_i maps the actions of all players to payoffs
- Mixed strategies result in one action; players calculate their expected payoff given the set of actions that might happen
- Written as tables with player actions on the axes, and the payoffs associated with each action profile in the boxes

Extensive Form Games:

- Pure strategies = an action at every information set
- u_i maps terminal histories to payoffs
- When converting to normal form, treat pure strategies as “actions”
 - Pure strategies go on the axes of the table
 - Pure strategy profiles correspond to histories which are associated with payoffs
- When considering best responses, rationalizable actions, strictly dominated actions, weakly dominated actions, etc., consider pure strategies (treat pure strategies as “actions”)

Notation

$\prod_{i=1}^N A_i = A_1 \times A_2 \times \dots \times A_N$	Taking one element (one action) from each set
$\{C, D\}$	A set of actions (e.g., a player’s pure strategies)
$\{(x, y, z), (a, b, c)\}$	A set of actions across players (e.g., the set of rationalizable actions)
(D, D)	A strategy profile (e.g., a Nash Equilibrium)
a_{-i}	An action by every player other than i
A_{-i}	The set of all combinations of actions by every player other than i
$\sigma_i = (p \circ C, (1 - p) \circ D)$	A mixed strategy for player i for some $p \in (0, 1)$

2. Equilibrium Concepts (Complete Information)

Nash Equilibrium (NE)

Rationalizability is sufficient to predict (D, D) in the Prisoner's Dilemma, because action C is never rationalizable for either player. But it does not let us predict anything for Battle of the Sexes, because both actions are rationalizable. Nash equilibria are a subset of rationalizable action profiles, further assuming each player has correct beliefs about what the other players will do. In Battle of the Sexes, the miscoordinated profiles are not NE because taking the action of the other player as given, it is optimal for me to match.

Strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_N^*)$ is a **Nash equilibrium (NE)** in the strategic game $\langle N, (A_i), (u_i) \rangle$ if for all i , $\sigma_i^* \in \Delta(A_i)$ and σ_i^* is a best response to the belief σ_{-i}^*

- Every player is best responding taking the strategies of the other players as given (no player has any unilateral incentive to deviate)
- Every player is playing a mixed strategy (assumed to be independent randomizations)

Every strategic game $\langle N, (A_i), (u_i) \rangle$ where all players have finitely many actions has a NE
(may include mixed strategies for one or more players)

- A strictly dominated strategy can never be a NE, so IESDS reduces the set under consideration
- In a finite strategic game, a mixed strategy is a NE if and only if for all i , every pure strategy in the support of σ_i^* is a best response to σ_{-i}^* (to mix, a player must be indifferent between the extremes)
- More generally, every game in which each action set is a convex compact subset of a Euclidian space and each payoff function is continuous has a NE

Finding Nash Equilibria: Continuous Action Spaces

✱ Taking as given the action of the other player, calculate the best response function

✱ A NE is any point where the best response functions intersect

✱ Player 1: $\max_{a_1} u_1(a_1, a_2)$ take $\frac{\partial u_1}{\partial a_1} = 0$ Player 2: $\max_{a_2} u_2(a_1, a_2)$ take $\frac{\partial u_2}{\partial a_2} = 0$

Solve the system of equations for a_1, a_2

✱ $B_1(a_2) = \arg \max_{a_1} u_1(a_1, a_2)$ (in NE, $B_1(a_2) = a_1$ and $B_2(a_1) = a_2$)

✱ Drawing a picture can be helpful

✱ Best response functions can be piecewise (e.g., to prevent negative quantities in firm competition)

✱ For best response correspondences, the set of NE is where the correspondences overlap

✱ Best response functions can also be used for discrete action spaces, as a function of p and q

Finding Nash Equilibria: Discrete Action Spaces

Do IESDS to eliminate non-rationalizable actions from consideration

Look for Pure NE

- * For each player, for each combination of actions by other players, underline his best response (e.g., for player 1, underline the biggest number in each column)
- * Any strategy profile with all underlined payoffs is a pure strategy NE

Look for Mixed NE

- * If there is more than one pure NE in a given row or column (a player has a choice between actions that all result in a NE), then any mix between those actions is also a NE
- * A player will only play a mixed strategy if he is indifferent to the actions being mixed
- * Player 1: Set $\mathbb{E}[X | q] = \mathbb{E}[Y | q]$ and solve for q (finds q that lets P1 mix between X and Y)
- * Setting the expected payoffs equal assumes that P1 is mixing between X and Y – also need to look for cases where P1 plays a pure strategy and P2 mixes

Games with finite NEs usually have an odd number (not always)

Write down clear strategy profiles e.g., $\sigma^* = (\sigma_1^*, \sigma_2^*)$ (not the payoffs)

Subgame Perfect Nash Equilibrium (SPNE)

Subgame perfect Nash equilibria are a subset of Nash equilibria for extensive form games. The Nash equilibrium concept sets no restrictions on player's strategies at information sets that are never reached on the equilibrium path (Nash equilibria can include non-credible threats). SPNE further assume that all players will play optimally at every information set of the game tree.

A **subgame** of an extensive game is a subset of the game such that:

- The subgame starts at some h that is the only history (node) in its information set
- All histories that follow from h are included in the subgame
- If any history (node) that is part of an information set is part of the subgame, then all histories (nodes) of this information set are part of the subgame

Notes about subgames:

- The number of subgames = the number of singleton information sets
- The history leading to (above) the subgame can include only part of an information set (e.g., a Prisoner's Dilemma followed by a Coordination Game)
- Proper subgames are subgames that are not the whole game itself
- Maximal subgames are subgames that have no further proper subgames

A Nash equilibrium σ^* is a **subgame perfect Nash equilibrium (SPNE)** in an extensive form game if it induces a Nash equilibrium in every subgame of this game

Every finite game has a SPNE (may include mixed strategies for one or more players)

Finding SPNE (Backward Induction)

- * Pick a NE in each maximal subgame
- * Replace the subgame with the payoffs from the NE
- * Repeat with the reduced game until the whole tree has been replaced
- * When pure NE are being selected, use arrows along actions to denote the SPNE path
- * If there are multiple NE at any point, then the process should be repeated for each (players can choose any of the NE, can be useful for inducing desired behavior further up in the tree)
- * For continuous action spaces, create a list of the actions taken and consider the optimal actions working backwards

Trembling Hand Perfect Equilibrium (THPE)

Trembling hand perfect equilibria are a subset of Nash equilibria for normal form games. THPE further assumes that there is a possibility of slight mistakes in a player's beliefs about what the other players will do (the equilibrium strategy must be robust to mistakes).

A **trembling hand perfect equilibrium (THPE)** of a finite strategic game is a mixed strategy profile σ^* with the property that there exists a sequence $\{\sigma^k\}_{k=0}^{\infty}$ of completely mixed strategy profiles that converges to σ^* such that for each player i the strategy σ_i^* is a best response to σ_{-i}^k for all values of k

- To show that a player's strategy can be part of a THPE, there only needs to be one sequence that it is robust to (not robust to every sequence)
- To show that a player's strategy cannot be part of a THPE, there must be an argument for why a sequence cannot be constructed
- Completely mixed NE are automatically THPE, as the sequence can be repetition of the NE itself (the definition does not capture robustness with respect to completely mixed strategy NE)

2 Player Games: σ^* is a NE in weakly undominated strategies $\Leftrightarrow \sigma^*$ is a THPE

3+ Player Games: σ^* is a THPE $\Rightarrow \sigma^*$ is a NE in weakly undominated strategies

Every finite strategic game has a THPE (may include mixed strategies for one or more players)

Finding THPE: 2 Players

- * Identify all of the NE of the game
- * Eliminate any NE where a player is using a weakly dominated strategy (do not do iterated elimination of any actions when evaluating this)

Showing σ^* is (or is not) a THPE

- * Confirm that σ^* is a NE – otherwise, it cannot be THPE
 - If asked to identify all THPE, start by identifying all NE of the game
 - If σ^* is the unique NE of the game, then it must be THPE
- * Confirm that no player is using a weakly dominated strategy – otherwise, it cannot be THPE (do not do iterated elimination of any actions when evaluating this)
- * Check if all players are using completely mixed strategies – if so, σ^* is a THPE by definition
- * For each player, let sequence p_i^k converge to σ_i^* such that every action has probability $\in (0, 1)$ along the sequence
 - Even strictly dominated strategies must have positive probability along the sequence
 - If player i has more than two actions, then p_i^k is a vector
Example: $(q_1^k, q_2^k, 1 - q_1^k - q_2^k) \rightarrow (1, 0, 0)$
- * Identify the conditions under which σ_i^* is a best response to $p_{-i}^k \forall k$
Example: $\mathbb{E}[\sigma_1^* | q^k] \geq \mathbb{E}[a_1 \neq \sigma_1 | q^k]$
- * Construct a sequence for each player such that the condition is met for all players (or show that the condition is met for any sequence)
- * To prove that σ is not THPE, assume that it is THPE towards a contradiction; implies there exists sequences meeting the criteria above; then show that for some player, σ_i can never be a best response to the sequences

ε -Equilibrium (ε -E)

The ε -equilibrium concept relaxes the assumption that all players are best responding (relaxes strict optimality). In an ε -E, one or more players can have a strictly positive but arbitrarily close to 0 incentive to deviate. ε -E is an expansion of NE in an attempt to better reflect behavior in actual games.

For any finite strategic game $\langle N, (A_i), (u_i) \rangle$, taking ε as given, σ^* is an **ε -equilibrium (ε -E)** if $u_i(\sigma^*) \geq u_i(\sigma'_i, \sigma_{-i}^*) - \varepsilon$ for all $\sigma'_i \in \Delta(A_i)$ and for all $i \in N$

- Every player could get a maximum of ε additional payoff by switching to any other strategy
- For a big enough ε , any strategy profile could be ε -E (ε is exogenous)
- A NE is always an ε -E; an ε -E is only an NE if the definition holds for $\varepsilon = 0$

Showing σ^* is (or is not) an ε -E

- * For each player i , identify the best response to σ_{-i}^* and the associated payoff
- * Identify the payoff associated with playing σ_i^*
- * Determine whether the best response payoff is within ε of playing σ_i^*

Correlated Equilibrium (CE)

We have assumed that players are using mixed strategies which are independent randomizations. The Correlated Equilibrium concept relaxes this assumption, allowing for players to jointly randomize over action profiles. For example, in Battle of the Sexes, the mixed strategy NE has each player independently randomizing, so with some probability the players will mismatch and both get 0. In a Correlated Equilibrium, the players could agree to play each NE 50% of the time, with no probability of ever playing the 0 payoff strategy profiles.

A **correlated equilibrium (CE)** of a strategic game $\langle N, (A_i), (u_i) \rangle$ consists of:

- A finite probability space with set of states Ω and probability measure π on Ω
- For each player $i \in N$, a partition \mathcal{P}_i of Ω (denotes which states the player can differentiate between)
- For each player $i \in N$, a function $\sigma_i : \Omega \rightarrow A_i$, with $\sigma_i(\omega) = \sigma_i(\omega')$ whenever $\omega \in P_i$ and $\omega' \in P_i$ for some $P_i \in \mathcal{P}_i$ (player i 's strategy is the same for every state within a set in his partition)

such that for every $i \in N$ and every function $\tau_i : \Omega \rightarrow A_i$ for which $\tau_i(\omega) = \tau_i(\omega')$ whenever $\omega \in P_i$ and $\omega' \in P_i$ for some $P_i \in \mathcal{P}_i$, we have $\sum_{\omega \in \Omega} \pi(\omega) u_i(\sigma_i(\omega), \sigma_{-i}(\omega)) \geq \sum_{\omega \in \Omega} \pi(\omega) u_i(\tau_i(\omega), \sigma_{-i}(\omega))$

Notes on CE:

- For each player, taking $\Omega, \pi(\omega)$, and \mathcal{P}_i as given, he has no reason to deviate to another strategy for any set in his partition
- Said another way, for every state $\omega \in \Omega$ that occurs with positive probability, the action $\sigma_i(\omega)$ is optimal given the other players' strategies and player i 's knowledge about ω

- We can think about Ω and π as a public correlation device (e.g., a coin or a third party)
- We can restrict the set of states to pure strategy profiles (every CE can be represented in this way)
- CE can result in individual or aggregate payoffs that are not possible with NE (mixed strategies) alone

For every mixed strategy NE σ^* of a finite strategic game $\langle N, (A_i), (u_i) \rangle$, there is a correlated equilibrium $\langle (\Omega, \pi), (\mathcal{P}_i), (\sigma_i) \rangle$ in which for each player $i \in N$ the distribution on A_i induced by σ_i is σ_i^*

CE Notation

$\Omega = \{x, y, z\}$ The set of states, each assigned to a pure strategy profile that is part of the CE

$\pi(x), \pi(y), \pi(z)$ The probability of each state happening

$\mathcal{P}_i = \{\{x\}, \{y, z\}\}$ For each player, a partition of the set of states

$\sigma_i = \begin{cases} T & \text{if } \{x\} \\ B & \text{if } \{y, z\} \end{cases}$ For each player, for each element in his partition, an associated pure strategy

Constructing a CE with Given Payoffs

Identify all pure NE of the game

Look for a mixed NE with given payoffs

Try to find a weighted average of pure NE

✱ Pick a set of pure NE to try weighting between

- Look first at weighting between 2 NE, then 3 NE
- May need to try a few combinations before finding one that works

✱ Assign each NE in the set an unknown probability and construct a system of equations

Example: NE x with p_x , NE y with p_y , NE z with p_z

$$\begin{aligned} p_x \cdot u_1(x) + p_y \cdot u_1(y) + p_z \cdot u_1(z) &= a && \text{(Player 1 expected payoff)} \\ p_x \cdot u_2(x) + p_y \cdot u_2(y) + p_z \cdot u_2(z) &= b && \text{(Player 2 expected payoff)} \\ p_x + p_y + p_z &= 1 && \text{(Probabilities must sum to 1)} \end{aligned}$$

✱ Solve the system of equations (e.g., with row operations on an augmented matrix)

Example:
$$\left[\begin{array}{ccc|c} u_1(x) & u_1(y) & u_1(z) & a \\ u_2(x) & u_2(y) & u_2(z) & b \\ 1 & 1 & 1 & 1 \end{array} \right]$$

✱ If there is a solution, then a CE has been found where each player's partition differentiates between all states; as the CE uses only pure NE, there is no need to check for incentives to deviate

If non-NE must be included:

✱ Pick a set of pure strategy profiles to try weighting between

✱ In a CE with the given payoffs in expectation:

- The probabilities across pure strategy profiles will sum to 1
- A player will always know what to play (each element in his partition will correspond to exactly one pure strategy)
- A player will never know for certain that he is not playing an optimal response (his partition will hide that information)
- For each player, for each element in his partition, there is no incentive to deviate from the CE strategy given that the other players follow the plan

✱ When a player does not know the state for sure, it is as though he is facing a mixed strategy

Example:

If Player 1 is in $\{x\}$, he knows Player 2 will play L and his best response is B (NE state)

If Player 1 is in $\{y, z\}$, then Player 2 is in $\{x, y\}$ and plays L with probability $\frac{p_y}{1 - p_x}$

Player 2 is in $\{z\}$ and plays R with probability $\frac{p_z}{1 - p_x}$

If we want Player 1 to play T in the second case, then $\mathbb{E}[T \mid \{y, z\}] \geq \mathbb{E}[B \mid \{y, z\}]$

Evolutionary Stable Strategies (ESS)

Evolutionary Stable Strategies are a refinement of NE to model situations in which the players' actions are determined by the forces of evolution. Consider a sequence of generations playing the same pairwise symmetric game. Strategies are inherited from the previous generation, with some small probability of mutations that lead to experimentation with other strategies. $u(\cdot)$ measures each organism's ability to survive, so strategies that lead to higher payoffs in the previous generation are used more often. Evolutionary Stable Strategies are what we expect to observe being used in the long run.

Let $(\{1, 2\}, (A, A), (u_i))$ be a strategic game where $u_1(a, b) = u_2(b, a) \forall a, b \in A$ (the game is symmetric). An **evolutionary stable strategy (ESS)** of the game is a mixed strategy σ^* for which (σ^*, σ^*) is a NE and $u(\sigma^*, \sigma) > u(\sigma, \sigma)$ for every best response σ to σ^* with $\sigma \neq \sigma^*$

- The intuition is to find strategies where no mutation would ever take hold in the population
- Typically assumed to be a finite game
- Every two-player symmetric strategic game in which each player has two pure strategies and the payoffs to the four strategy profiles are different has an ESS (other two-player symmetric games may not)

Finding ESS

- * Find all symmetric NE of the game (pure and mixed)
- * For each NE σ^* , identify the best responses to σ^*
- * If σ^* is the only best response to σ^* , then σ^* is an ESS
- * Otherwise, for each best response $\sigma \neq \sigma^*$, confirm that $u(\sigma, \sigma)$ (mutant vs. mutant payoff) is strictly less than $u(\sigma^*, \sigma)$ (non-mutant vs. mutant payoff)

3. Bayesian Games

A game has **perfect information** if all information sets are singletons

A game has **imperfect information** if at least one information set has multiple nodes

A game has **incomplete information** if at least one player does not know the payoff that some player receives from some strategy profile (or terminal node)

We have assumed until now that all players know each other's preferences over all strategy profiles (games of complete information). We want to expand to environments or situations where we are uncertain about the motives or incentives of the people we are acting (games of incomplete information). We transform these games into extensive games with complete but imperfect information.

A **Bayesian game** $\langle N, (A_i), (u_i), (\Theta_i), F \rangle$ consists of:

- A finite set of players N
- For each player $i \in N$, a non-empty set of actions A_i
- For each player $i \in N$, a utility function $u_i : \prod_{i=1}^N A_i \times \prod_{i=1}^N \Theta_i \rightarrow \mathbb{R}$ ($u_i : A \times \Theta \rightarrow \mathbb{R}$)
 - Utility is determined by the actions of every player and the types of every player
 - A player considers his payoffs as an expectation of the types he might face
- For each player $i \in N$, a set of types Θ_i
 - At the beginning of the game, Nature chooses realizations of each random variable Θ_i
 - Each player observes only his own type $\theta_i \in \Theta_i$ but not that of the others
- Prior probability distribution over type profiles $F : \Theta \rightarrow [0, 1]$
 - F is common knowledge
 - Typically we assume that the realization of each type is independent (e.g., we consider $\pi_i(H) = \pi_i(L) = \frac{1}{2}$)

A **pure strategy for player i in a Bayesian game (pure decision rule)** is a mapping $s_i : \Theta_i \rightarrow A_i$ where $s_i(\theta_i) \in A_i$ is a pure strategy of the basic game chosen by type θ_i

- The decision rule specifies what action a player will take for each of her possible type realizations
- We can similarly consider mixed decision rules
- Given a profile of decision rules for all players, player i 's expected utility is:

$$\tilde{u}_i(s_1, \dots, s_N) = \mathbb{E}[u_i(s_1(\theta_1), \dots, s_N(\theta_N), \theta_i, \theta_{-i})]$$

A **Bayesian Nash equilibrium (BNE)** of the Bayesian game $\langle N, (A_i), (u_i), (\Theta_i), F \rangle$ is a profile of decision rules (s_1, \dots, s_N) such that $\tilde{u}_i(s_i, s_{-i}) \geq \tilde{u}_i(s'_i, s_{-i})$ for all i and all $s'_i : \Theta_i \rightarrow A_i$

- The expected utility for each player of following his decision rule is at least as great as any other decision rule, given the decision rules of all the other players

A profile of decision rules s is a BNE if and only if, for all i , for all $\theta_i \in \Theta_i$ occurring with positive probability, and for all $\tilde{s}_i \in A_i$: $\mathbb{E}_{\theta_{-i}}[u_i(s_i(\theta_i), s_{-i}(\theta_{-i}, \theta_i, \theta_{-i}) \mid \theta_i)] \geq \mathbb{E}_{\theta_{-i}}[u_i(\tilde{s}_i(\theta_i), s_{-i}(\theta_{-i}, \theta_i, \theta_{-i}) \mid \theta_i)]$

- We can find an optimal response for a player (to the decision rules of all the other players) one of his types at a time (treat each type as a separate player)

Every Bayesian game $\langle N, (A_i), (u_i), (\Theta_i), F \rangle$ where all players have finitely many actions and finitely many types has a (possibly mixed) BNE

- We will typically focus on pure, symmetric BNE

Finding BNE: Discrete Types, Discrete Action Spaces

- * For each player, identify how many types he has - his strategy must include this many components
- * Consider all possible combinations of types facing each other and write out the potential games
- * For Player 1, construct a “short & fat” table of expected payoffs for each type (pure strategies for P1 type vs. all possible P2 strategies)
- * Construct an overall table of expected payoffs for Player 1 (all possible P1 strategies vs. all possible P2 strategies)
- * For Player 2, construct a “tall & skinny” table of expected payoff for each type (all possible P1 strategies vs. pure strategies for P2 type)
- * Construct an overall table of expected payoffs for Player 2
- * Combine tables into normal form game and find the (pure or mixed) BNE
- * Example: $\pi_1(S) = p$, $\pi_1(W) = 1 - p$, $\pi_2(S) = q$, $\pi_2(W) = 1 - q$

Each player will have two components to his strategy, e.g., $A_S N_W$

4 potential games: $(S, S), (S, W), (W, S), (W, W)$ (S, W) game considers $\{A_S, N_S\}$ vs. $\{A_W, N_W\}$

P1=S “short & fat” has $\{A_S, N_S\}$ vs. $\{A_S A_W, A_S N_W, N_S N_W, N_S N_W\}$ with payoffs weighted q

P1 overall payoffs has $\{A_S A_W, A_S N_W, N_S N_W, N_S N_W\}$ vs. $\{A_S A_W, A_S N_W, N_S N_W, N_S N_W\}$, and is weighting P1=S vs. P1=W by p

P2=S “tall & skinny” has $\{A_S A_W, A_S N_W, N_S N_W, N_S N_W\}$ vs. $\{A_S, N_S\}$ weighted by p

P2 overall payoffs is weighting P2=S vs. P2=W by q

Finding BNE: Discrete Types, Continuous Action Spaces

- * For each player, identify how many types he has - his strategy must include this many components
- * Calculate best response functions for each player type, given the decision rule of the other players and the probability distribution over their types
- * Example:

$$\text{Player 1, Type H: } \max_{a_1^H} u_1^H(a_1^H, a_2) \quad \text{take } \frac{\partial u_1^H}{\partial a_1^H} = 0$$

$$\text{Player 1, Type L: } \max_{a_1^L} u_1^L(a_1^L, a_2) \quad \text{take } \frac{\partial u_1^L}{\partial a_1^L} = 0$$

$$\text{Player 2: } \max_{a_2} u_2((\pi(H)a_1^H + \pi(L)a_1^L), a_2) \quad \text{take } \frac{\partial u_2}{\partial a_2} = 0$$

Solve the system of equations for a_1^H, a_1^L, a_2

Finding BNE: Continuous Types

- * Find each player's best response function as a function of his type
 - For discrete action spaces, this is a threshold (e.g., accept if $r(\theta) \leq w$, reject if $r(\theta) > w$)

Finding BNE: Auctions
(Continuous Types, Continuous Action Spaces)

- * Assume a strictly increasing, bidding function of a given form
 - Typically linear: $b(s) = \alpha s + c$
 - Can also be quadratic (all-pay), log, etc.
- * There is a distribution of signals (e.g., $s_i \sim U[0, 1]$)
- * Solve the problem from Player i 's perspective
- * $s(b)$ = the signal that induces Player i to bid b
- * Assume opponents are all using the same bidding function and define $\bar{s}(b)$ = the signal that induces opponents to bid b
- * Choose $b \mid s_i$ to maximize $P(\text{win} \mid b) \cdot \mathbb{E}[\text{payoff} \mid \text{win}] + P(\text{lose} \mid b) \cdot \mathbb{E}[\text{payoff} \mid \text{lose}]$
- * $F(\bar{s}(b))$ = the probability one opponent's signal is less than the signal that induces him to bid b
 - For $s_i \sim U[0, 1]$, $F(\bar{s}(b)) = \bar{s}(b)$
- * $\max_{b \mid s_i} = \{F(\bar{s}(b))^{n-1} \cdot \mathbb{E}[\text{payoff} \mid \text{win}] + (1 - F(\bar{s}(b))^{n-1}) \cdot \mathbb{E}[\text{payoff} \mid \text{lose}]\}$
- * In all-pay, $\mathbb{E}[\text{payoff} \mid \text{lose}] = b$; otherwise it generally equals 0 so second term goes away
- * In private value, $\mathbb{E}[\text{payoff} \mid \text{win}] = s_i - b$
- * In common value, $\mathbb{E}[\text{payoff} \mid \text{win}] = \mathbb{E}[v \mid s_i, s(b) > s_j, \forall j \neq i] - b$
 - $= \frac{1}{n} \sum s_j - b = \frac{1}{n} [s_i + (n-1)\mathbb{E}[s_j \mid \bar{s}(b)]] - b = \frac{1}{n} \left[s_i + (n-1) \frac{\bar{s}(b)}{2} \right] - b$
 - Because the bidding function is strictly increasing, my bid $>$ opponent bid is the same as my signal $>$ opponent signal
- * Take the derivative w.r.t. b of max equation and set $= 0$
- * Replace $\bar{s}'(b)$ with derivative of bidding function form (e.g., linear, $\bar{s}'(b) = \frac{1}{\alpha}$)
- * Replace $\bar{s}(b) = s_i = s$ (now assume opponents and Player i have the same bidding function)
- * Solve for b as a function of s and simplify

Global Games

In a game where the underlying state of the world is common knowledge, there may be multiple NE. Even if there is a Pareto-dominant NE, our existing solution concepts do not let us predict how often this NE is played. Global games are games of incomplete information where players receive correlated signals of the underlying state of the world. We use global games to make predictions about what NE would be played as the range of signals goes to zero.

- Some variable $\theta \sim U[\underline{\theta}, \bar{\theta}]$ is observed with uniform noise on $[\theta - \varepsilon, \theta + \varepsilon]$; noise is iid across players
- x_i is the signal received by player i
- Because of the uniform distribution, player i knows that the true θ is distributed uniformly within ε of his signal; $\theta \sim U[x_i - \varepsilon, x_i + \varepsilon]$ and $\mathbb{E}[\theta | x_i] = x_i$
- k is the threshold signal; we assume a symmetric strategy across all players, where one strategy is followed when $x_i \leq k$ and the other strategy is followed when $x_i > k$ (k is generally a function of ε)
- We focus on cases where k is sufficiently close to the middle of $[\underline{\theta}, \bar{\theta}]$ and ε is sufficiently small to not worry about weird edges of distributions
- At $x_i = k$, $\mathbb{E}[u_i(C) | x_i] = \mathbb{E}[u_i(D) | x_i]$ (plug in k for x_i at the very end)
- $\mathbb{E}[u_1(D) | x_1] = \int_{x_1 - \varepsilon}^{x_1 + \varepsilon} [u_1(D, C) \cdot P(x_2 > k | \theta) + u_1(D, D) \cdot P(x_2 \leq k | \theta)] f(\theta) d\theta$
 - The expected value to Player 1 integrated over possible values of θ , assuming Player 2 also follows the threshold strategy
 - $P(x_2 \leq k | \theta) = \frac{k - \theta + \varepsilon}{2\varepsilon}$
 - $f(\theta) = \frac{1}{2\varepsilon}$
- Solve for k as a function of ε and find $\lim_{\varepsilon \rightarrow 0} k(\varepsilon)$
- The converged k gives our prediction of which NE will be played for different values of true θ

4. Perfect Bayesian and Sequential Equilibria

SPNE are a refinement of NE that assumes players will play optimally at every node of the game tree (sets restrictions on strategies at nodes off of the equilibrium path). In extensive form games with imperfect information, a similar issue can arise. In particular, a NE may include a non-optimal action at an information set that is never reached. In games of imperfect information, the structure of the game may limit the number of subgames (limiting the relevance of SPNE).

Perfect Bayesian Equilibria and Sequential Equilibria set restrictions on how players act at information sets that are off of the equilibrium path. They do this by considering an equilibrium both a strategy profile and a belief system (NE implicitly include beliefs that are always correct).

Perfect Bayesian Equilibria and Sequential Equilibria each include:

- A strategy profile σ that prescribes actions (or probability distributions over actions) at every information set
- A belief system μ that assigns beliefs (or probability distributions) over nodes at every information set
 - When information sets are singletons, beliefs are trivial (the player believes he is at the node with probability 1); only write beliefs for information sets with multiple nodes
 - $\mu_1(a)$ is Player 1's belief that she is at node a conditional upon reaching that information set

Perfect Bayesian Equilibrium (PBE)

PBE are a subset of NE or BNE (whichever applies), where off-path actions (or probability distributions over actions) are justified by some (any) belief.

A strategy profile σ is **sequentially rational** given beliefs μ if at every information set, given the beliefs prescribed by μ , no individual wants to deviate from the piece of the strategy assigned to her there

- Taking beliefs as given, individuals are best-responding at every information set

(σ, μ) is a **perfect Bayesian equilibrium (PBE)** if σ is sequentially rational given μ and each agent updates beliefs μ using Bayes' Rule whenever it applies (i.e., at every information set on the equilibrium path)

- On-path, beliefs must be updated with Bayes' Rule (equivalent to correct beliefs)
- Off-path, beliefs can be anything

Finding PBE

- * Identify all NE or BNE
- * For each NE or BNE σ , check all on-path and off-path information sets:
- * On-Path:
 - For each multi-node information set, specify the associated **belief** using Bayes' Rule (equivalent to correct beliefs)
 - Since σ is a NE or BNE, we know that it is optimal on-path given correct beliefs
 - Otherwise, must also check that σ is optimal for all players on-path given correct beliefs
- * Off-Path:
 - For each information set, find any associated **belief** such that σ is **optimal**
 - For singleton information sets, beliefs are trivial – but optimality must still be checked!
- * Write down the strategy σ and the belief system μ that support the PBE
(alternatively, the range of beliefs that support σ as part of a PBE)

Sequential Equilibria (SE)

SE are a subset of PBE, where off-path actions (or probability distributions over actions) are justified by consistent (“reasonable”) beliefs. They are also a subset of SPNE.

A system of beliefs μ is **consistent** with strategy profile σ if there exists a sequence (σ^m, μ^m) of strategy profiles and beliefs such that:

- (i) $(\sigma^m, \mu^m) \rightarrow (\sigma, \mu)$ as $m \rightarrow \infty$
- (ii) σ^m is completely mixed for every m (assigns strictly positive probability to every action at every information set)
- (iii) For each m , μ^m is derived from σ^m by applying Bayes' Rule to every information set

(σ, μ) is a **sequential equilibrium (SE)** if σ is sequentially rational given μ and μ is consistent given σ

- On-path, beliefs must be updated with Bayes' Rule (equivalent to correct beliefs)
- Off-path, beliefs must be “reasonable” ($\mu^m \rightarrow \mu$ where μ^m is updated with Bayes' Rule)
- σ^m does not need to be best-responding to μ^m along the sequence (different from THPE)

Every finite extensive form game with imperfect information has a (possibly mixed) SE

Finding SE

- * Identify all PBE or SPNE (use SPNE if there are proper subgames)
- * For each PBE or SPNE σ , check whether every information set is reached with positive probability (if so, then σ is an SE)
- * Otherwise, check all off-path information sets
 - Intuitively, the belief must reflect that if an off-path information set is reached, it is more likely that one mistake was made rather than two
 - No need to verify sequences if not explicitly asked for and are sure that (σ, μ) is a SE
- * Write a completely mixed σ^m that converges to σ
 - Focus on information sets before the information set in question
 - Write σ^m using $\varepsilon^m \rightarrow 0$, etc.
 - For more actions, try $\sigma_1^m = ((1 - \varepsilon^m - (\varepsilon^m)^2) \circ a, \varepsilon^m \circ b, (\varepsilon^m)^2 \circ c)$
- * Derive μ^m for the information set using Bayes' Rule
 - $$\mu_i^m(a) = \frac{\text{Probability Player } i \text{ is at node } a}{\text{Probability at node } a + \text{Probability at other nodes in the information set}}$$
- * Try to find $\lim_{m \rightarrow \infty} \mu^m = \mu$
- * If μ equals some number $[0, 1]$ for all sequences ε^m , then check that σ is optimal given μ
- * If μ is ambiguous (depends on the exact sequences), then construct any sequences such that $\mu^m \rightarrow \mu$ that makes σ optimal
 - There is no information on which the player in question can update (e.g., they don't know which player made a mistake, or given that a mistake happened they don't know which other action was taken)
 - μ^m can converge anywhere in $[0, 1]$
 - Identify the range of μ such that σ is optimal and converge there (only need one sequence to show σ is part of an SE)
 - $\varepsilon^m = \frac{1}{m}$ is a good place to start
- * Write down the strategy σ and the belief system μ that support the SE (alternatively, the range of beliefs that support σ as part of a SE)
 - Include on-path beliefs as well as these are technically part of the SE

- * Example for Player 2's information set: $\sigma = (AD, X)$ where A ends the game

$$\sigma_1^m = (((1 - \varepsilon^m) \circ A, \varepsilon^m \circ B), (\gamma^m \circ C, (1 - \gamma^m) \circ D))$$

$$\mu_2^m(C) = P(\text{at } C \mid \text{P2's information set is reached}) = \frac{\varepsilon^m \cdot \gamma^m}{\varepsilon^m \cdot \gamma^m + \varepsilon^m(1 - \gamma^m)} = \frac{\varepsilon^m \cdot \gamma^m}{\varepsilon} = \gamma^m$$

$$\mu_2^m(C) \rightarrow 0 = \mu_2(C)$$

In order for σ to be an SE, X must be optimal given $\mu_2(C) = 0$

Signaling Games

Signaling games are a subset of extensive form games with imperfect information where PBEs are used. Player 1 knows his type and uses his action to send a signal that Player 2 observes and can act on.

- **Stage 0:** Nature chooses types $\theta \in \Theta$ of player 1 from distribution p
- **Stage 1:** Player 1 observes θ and chooses $a_1 \in A_1$
- **Stage 2:** Player 2 observes a_1 and chooses $a_2 \in A_2$
- **Payoffs:** $u_1(a_1, a_2, \theta), u_2(a_1, a_2, \theta)$

A **PBE in a signaling game** is a strategy profile $s_1(\theta), s_2(a_1)$, and beliefs $\mu_2(\theta | a_1)$ such that:

- (i) **Player 1's strategy is optimal given Player 2's strategy**

For all $\theta \in \Theta$, $s_1(\theta)$ solves $\max_{a_1 \in A_1} u_1(a_1, s_2(a_1), \theta)$

- (ii) **Player 2's beliefs are compatible with Bayes' Rule**

If any type of Player 1 plays a_1 with positive probability, $\mu_2(\theta | a_1) = \frac{P(s_1(\theta) = a_1)p(\theta)}{\sum_{\theta' \in \Theta} P(s_1(\theta') = a_1)p(\theta')}$

- (iii) **Player 2's strategy is optimal given his beliefs and given Player 1's action**

For all $a_1 \in A_2$, $s_2(a_1)$ solves $\max_{a_2 \in A_2} \sum_{\theta \in \Theta} u_2(a_1, a_2, \theta)\mu_2(\theta | a_1)$

PBE in signaling games fall into three categories:

- 1 **Separating:** Each type of Player 1 uses a different strategy. Hence, by observing a_1 , Player 2 can always back out θ
- 2 **Pooling:** All types of Player 1 use the same strategy. Hence, from observing a_1 , Player 2 cannot infer anything about θ
- 3 **Semi-Separating:** Some types of Player 1 use the same strategy, while some types use different strategies. Hence there is some, but not perfect, learning about θ from a_1

We typically do not worry about SE in signaling games. Either every information set is reached with positive probability (separating and semi-separating), or there is no information on which to update and beliefs can be anything (Player 2 doesn't know which type of Player 1 made a mistake)

Signaling PBE: Discrete Action Spaces

Separating PBE:

- * Each type picks a different strategy
- * Assuming the number of types matches the number of strategies, every information set is reached with positive probability so beliefs must be correct
- * Check that no Player 1 given type and no Player 2 given signal has an incentive to deviate

Pooling PBE:

- * All types pick the same strategy
- * Try pooling on one strategy first, then another, etc.
- * Beliefs at Player 2's on-path information set must reflect the underlying distribution of types
- * Check that given on-path beliefs, Player 2 is best responding
- * Choose Player 2's off-path strategy such that Player 1 does not have an incentive to deviate
- * Construct any belief for Player 2's off-path information set such that his off-path strategy is optimal

Semi-Separating:

- * At least one type is mixing between strategies and at least one type is not mixing

Write down σ and μ

Signaling PBE: Continuous Action Spaces

Player 2 has some belief function $\mu(a_1) \in [0, 1]$

Player 2 has some action function $\mu \rightarrow A_2$

Separating PBE:

- * $\mu_Y(a_{1X}) = 0$ and $\mu_Y(a_{1Y}) = 1$ (for some types X and Y), and Player 2 will map these beliefs into distinct actions a_{2X} and a_{2Y}
- * Given a_{2X} and a_{2Y} , each type of Player 1 will choose a_1
 - For the Player 1 “low type”, the worst belief Player 2 can have is already occurring (beliefs can’t be used to incent the PBE); this type must be optimizing given they are perfectly identified
 - $u_X(a_{1X}) \geq u_X(a_{1Y})$ (type X must not deviate to type Y)
 - $u_Y(a_{1Y}) \geq u_Y(a_{1X})$ (type Y must not deviate to type X)
- * Off-path, Player 2’s beliefs (and thus actions) must be such that no type of Player 1 has an incentive to deviate to any $a_1 \notin \{a_{1X}, a_{1Y}\}$ (this can be “low” belief / action for everything off-path or more general)

Pooling PBE:

- * $a_{1X} = a_{2X} = a_1^*$
- * $\mu(a_1^*)$ must be equal to the prior distribution of types, and Player 2 maps this belief into a_2^*
- * Off-path, Player 2’s beliefs (and thus actions) must be such that no type of Player 1 has an incentive to deviate to any $a_1 \neq a_1^*$ (this can be “low belief / action for everything off-path or more general)
- * To check for deviations, identify the most advantageous deviation and the Player 1 type that is more likely to deviate (has lower utility in the PBE, is the “low” type with nothing to lose, etc.)

Semi-Separating:

- * Player 2 will observe two different actions; one that corresponds to a for-sure belief (e.g., $\mu = 0$) and one that corresponds to a Bayes’ Rule updated belief $\mu \notin \{0, 1\}$
- * The Player 1 type that is mixing must be indifferent to the two actions, while the other type must strictly prefer one action
- * Off-path, Player 2’s beliefs (and thus actions) must be such that no type of Player 1 has an incentive to deviate (this can be “low belief / action for everything off-path or more general)

Write down σ and μ

5. Infinitely Repeated Games

We have until now only considered games over finite time horizons. Infinitely repeated games are a subset of dynamic games where the same strategic interaction occurs over infinite time horizons. They may represent situations where the interaction actually extends forever, or where players know the game is finite but have no knowledge of when it will end. Infinitely repeated games can support a much richer set of possible equilibria, especially when actions are strategically linked over time.

A **stage game** is a one-period simultaneous move (normal-form) game of complete information given by $\langle N, (A_i), (v_i) \rangle$ where A is finite.

A **infinitely repeated game** is a game formed by a repetition of the stage game for $T = \infty$ periods. We assume perfect monitoring such that at the end of every period, every player observes the actions chosen by all players in that period.

A **strategy in a repeated game** is $\sigma_i : \mathcal{H} \rightarrow \Delta(A_i)$ where \mathcal{H} is the space of all possible histories in the repeated game. It specifies what action (or probability distribution over actions) to take in any period following any history at that period. In an infinitely repeated game, clearly \mathcal{H} is also infinite; strategies thus identify “states” of histories and map these states to actions.

Payoffs are written as infinite sequences of expected payoffs: $u_i(\{v_i^t\}_{t=1}^{\infty}) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} v_i^t$

- We assume exponential discount factor $\delta \in [0, 1)$
- The sequence of payoffs is normalized by $(1 - \delta)$ for convenience
- $v_i^t = v^* \quad \forall t \Rightarrow u_i(\{v_i^t\}_{t=1}^{\infty}) = v^*$
- $u_i(\{v_i^t\}_{t=1}^{\infty}) = (1 - \delta)v_i^1 + \delta u_i(\{v_i^t\}_{t=2}^{\infty})$

A **profitable one-shot deviation (POSD)** for player i is a strategy $\sigma'_i \neq \sigma_i$ such that (a) there is a unique history h^t such that for all $\tilde{h}^t \neq h^t$, $\sigma'_i(\tilde{h}^t) = \sigma_i(\tilde{h}^t)$ and (b) $u_i(\sigma'_i, \sigma_{-i} \mid h^t) > u_i(\sigma \mid h^t)$

- There is exactly one history where σ'_i deviates from σ_i , and the deviation is profitable conditional on reaching that history

One-Shot Deviation Principle:

- A strategy profile σ is a NE if and only if there are no POSD on the equilibrium path
- A strategy profile σ is a SPNE if and only if there are no POSD anywhere
- Intuitively, if there exists a profitable deviation then there exists a POSD
- These results are powerful because they limit the deviations we have to check from infinity to the smallest possible deviations
- These results only hold because $\delta < 1$

If the stage game has a NE, then the infinitely repeated game has a SPNE (repeat the NE always)

Checking if σ is NE / SPNE in Infinitely Repeated Game

Given $\sigma_i : \mathcal{H} \rightarrow \Delta(A_i)$, identify:

- * The set of states Ω such that $\Omega \rightarrow \Delta(A_i)$
 - Each $\omega \in \Omega$ maps to a unique action
 - ω_{CC} represents the state where both players are supposed to play C in the next stage game
- * The transition function such that $\Omega \times \prod_{i=1}^n A_i \rightarrow \Omega$
 - For each state, every action profile will create a transition to (another or the same) state
- * Which state the game begins in (ω_0 “transitions” immediately into a state)
- * Diagramming the states and transitions with circles and arrows can be helpful

Check for POSD in On-Path States (NE):

- For every player, check for POSD in on-path states
- Compare the value of starting in the state and following the strategy $V_i(\omega)$, to the value of starting in the state, deviating once, and then following the strategy for the rest of time $\tilde{V}_i(\omega)$
- Use the diagram to identify the future sequence of states
- If $V_i(\omega) - \tilde{V}_i(\omega) \geq 0$ for all i and all on-path ω , then σ is a NE
- Solve for the range of δ that makes σ an NE (use the most restrictive criteria if multiple)
- If a state maps to playing a NE of the stage game, there is no POSD
- Example:

$$V_1(\omega_{CC}) = u_1(C, C) + \delta V_1(\omega_{CC}) = (1 - \delta) [5 + \delta 5 + \delta^2 5 \dots]$$

$$\tilde{V}_1(\omega_{CC}) = u_1(D, C) + \delta V_1(\omega_{DD}) = (1 - \delta) [6 + \delta 2 + \delta^2 2 \dots]$$

Check for POSD in Off-Path States (SPNE):

- For every player, check for POSD in every off-path state (as above)

Some strategies for infinitely repeated Prisoner’s Dilemma: (must consider both players’ strategies):

- **Always Defect:** Both players following this strategy is a SPNE (repeating the stage NE always)
- **Always Cooperate:** Both players following this strategy is not a NE
- **Grim Trigger:** In the first round, play C . In subsequent rounds, play C if the history is such that neither player has ever played D ; otherwise play D . For a range of δ (if players are sufficiently patient), both players following Grim Trigger is a SPNE.
- **Tit-for-Tat:** In the first round, play C . In subsequent rounds, play whatever the opponent played in the previous period. Both players following Tit-for-Tat is a NE for a range of δ ; it is an SPNE for at most one value of δ .
- If both players use Always Cooperate, Grim Trigger, or Tit-for-Tat, we will observe (C, C) forever, but they are very different strategies

Folk Theorems

Folk theorems specify what payoff profiles can be achieved by a SPNE of an infinitely repeated game.

The **set of feasible payoffs**, \mathbf{V} , is defined as the convex hull of the set of all payoff vectors that are attained by some action profile of the stage game

- $V = \text{co}(\tilde{V})$ where $\tilde{V} = \{(v_1(a), \dots, v_N(a))\}_{a \in A}$
- A convex hull of \tilde{V} is the smallest convex set that contains \tilde{V} and all convex combinations of \tilde{V}
- For 2 players, graph the payoffs of all possible action profiles and connect the points with straight lines

Player i 's **Nash-threat** payoff is $\underline{v}_i = \inf\{v_i \mid \exists \text{ stage-game (possibly mixed) NE } \sigma \ni v_i(\sigma) = v_i\}$

- The lowest payoff player i can receive in a pure or mixed NE

Nash-Threat Folk Theorem: Given a stage game, pick any $v \in V$ such that for all i , $v_i > \underline{v}_i$. There exists a $\underline{\delta} \in [0, 1)$ such that for all $\delta > \underline{\delta}$, there is a SPNE of the infinitely repeated game with average discounted payoff profile v

- If players are patient enough, any payoff where all players get more than their Nash-threat payoff can be achieved in an SPNE
- Intuitively, can incent a player to follow the strategy so long as his payoff on-path is higher than getting his Nash-threat payoff forever
- For 2 players, find Nash-threats $(\underline{v}_1, \underline{v}_2)$. Anything in V from 12:00 to 3:00 can be achieved in SPNE

Player i 's **minmax** payoff is given by $v_i^0 = \min_{\sigma_{-i}} \max_{\sigma_i} v_i(\sigma_i, \sigma_{-i})$ and a vector of payoffs v is **strictly**

individually rational if for all i , $v_i > v_i^0$

- v_i^0 is the largest payoff that Player i can guarantee himself, even if all of the other players coordinate against him – check all pure strategies of other players and find the minimum best response

Minmax Folk Theorem: Given a finite stage game, let $V^* \subset V$ be the subset of feasible and strictly individually rational payoffs. If v is in the interior of V^* , then there is a $\underline{\delta} \in [0, 1)$ such that for all $\delta > \underline{\delta}$, there is a SPNE of the infinitely repeated stage game with average discounted payoff profile v

Constructing SPNE in Infinitely Repeated Game

- * Start by identifying NE of the stage game; no players will ever have a POSD from a NE
- * Easiest construction is a Nash Reversion strategy; if players don't follow the strategy, the punishment is getting a bad NE forever
- * Otherwise try 1- or 2-period "punishment" states (punishments for deviations don't have to be symmetric across players)
- * SPNE in infinitely repeated games can include public coordination devices, different strategies for even vs. odd rounds, etc.

Appendix

Relating Solution Concepts

Rationalizable Actions: The broadest solution concept. For each player, there is a set of actions that is rationalizable. Rationalizability allows for correlated strategy beliefs, but does not require correct beliefs.

Nash Equilibria (NE): A refinement of rationalizable actions, further assuming each player has correct beliefs about what the other player will do. NE are mixed strategy profiles, made up of rationalizable actions by all players, where no player has an incentive to deviate. NE and its refinements do not allow for correlated strategies.

Subgame Perfect Nash Equilibria (SPNE): A subset of NE for extensive form games. SPNE assumes that all players will play optimally at every information set in a game.

Trembling Hand Perfect Equilibria (THPE): A subset of NE for normal form games. THPE assumes that there is a possibility of slight mistakes in a player's beliefs about what the other players will do.

SPNE and THPE: Are two different solution concepts. The THPE of a game may be a subset of the SPNE, or the SPNE may be a subset of the THPE.

ϵ -Equilibria (ϵ -E): ϵ -equilibrium relaxes the best response assumption, such that one or more players can have a strictly positive but arbitrarily close to 0 incentive to deviate. ϵ -E is an expansion of NE in an attempt to better reflect behavior in actual games.

Correlated Equilibrium (CE): NE and its refinements use mixed strategy profiles, which we have assumed are independent randomizations. CE allows for players to use strategies such that they jointly randomize over action profiles. CE requires correct beliefs, but allows correlated strategies. A NE (pure or mixed) is always a CE.

Evolutionary Stable Strategies: A subset of NE when modeling situations in which the players' actions are determined by the forces of evolution.

Bayesian Nash Equilibrium (BNE): The equivalent of NE in games of incomplete information that have been transformed into games with complete but imperfect information. Similar to NE, BNE do not place any restrictions on strategies off of the equilibrium path (can include things like non-credible threats).

Mixed NE and BNE: Consider a sequence of Bayesian games that converges to a complete information normal form game (e.g., two players have different, randomized payoffs but the range of randomizations is approaching zero). In general, every NE (pure or mixed) of the complete information is the limit of NE for the sequence of Bayesian games. As players have a unique best response along the sequence, they are not

randomizing. Thus mixed NE in complete information games can be viewed as pure strategy equilibria in a close-by Bayesian game.

Global Games: Games of incomplete information where players receive correlated signals of the underlying state of the world. We use global games to make predictions about what NE would be played as the range of signals goes to zero.

Perfect Bayesian Equilibrium (PBE): In games of incomplete information, a subset of NE or BNE (whichever applies) where off-path strategies are justified by some (any) belief.

Sequential Equilibrium (SE): A subset of PBE, where off-path actions (or probability distributions over actions) are justified by consistent (“reasonable”) beliefs. They are also a subset of SPNE.

NE, PBE, and SE as Long-Term Play: We can think of NE as an equilibrium reached after long-term play, where players learn about each others’ strategies. Similarly, we can think of PBE as where players plan to behave optimally off-path, but have never observed what happens off-path so can believe anything. SE are where the other players make shaky-hand mistakes every so often, so each player learns what happens off-path and plan to behave optimally in response.

Signaling Games: A subset of extensive form games with imperfect information where PBEs are used. Player 1 knows his type and uses his action to send a signal that Player 2 observes and can act on.

Infinitely Repeated Games: A subset of dynamic games where the same strategic interaction occurs over infinite time horizons. We look at NE and SPNE in infinitely repeated games.

Uniform Distribution

	$\theta \sim U[a, b]$	$\theta \sim U[0, 1]$
$f(\theta^*) = P(\theta = \theta^*) :$	$\frac{1}{b - a}$	1
$F(\theta^*) = P(\theta \leq \theta^*) :$	$\frac{\theta^* - a}{b - a}$	θ^*
Mean:	$\frac{a + b}{2}$	$\frac{1}{2}$
Variance:	$\frac{(b - a)^2}{12}$	$\frac{1}{12}$

Geometric Series

$$(1 - \delta^2) = (1 + \delta)(1 - \delta)$$

$$(1 - \delta^3) = (-1)(\delta - 1)(\delta^2 + \delta + 1)$$

$$\sum_{t=1}^T \delta^{t-1} a = a + \delta a + \delta^2 a + \delta^3 a \dots \delta^{T-1} a = \frac{a - a \cdot \delta^T}{1 - \delta}$$

$$\sum_{t=1}^{\infty} \delta^{t-1} a = a + \delta a + \delta^2 a + \delta^3 a \dots = \frac{a}{1 - \delta} \quad (\text{in the limit!})$$

$$\sum_{t=1}^{\infty} (\delta^2)^{t-1} a = a + (\delta^2) a + (\delta^2)^2 a + (\delta^2)^3 a \dots = a + \delta^2 a + \delta^4 a + \delta^6 a = \frac{a}{1 - \delta^2}$$