Game Theory

Topic 1: Knowledge

Formalizing Knowledge

Ω: possible states of the world (finite)

$$www - wwb - wbw - wbb - bww - bwb - bbw - bbb$$
$$\Omega = \{1,2,3,4,5,6,7,8\}$$

$E \subseteq \Omega$: events

$$E = \{2,6\}\text{: either } wwb \text{ or } bwb$$

H_i, a partition of Ω: what a person can distinguish

$$H_1 = \{\{1,5\},\{2,6\},\{3,7\},\{4,8\}\}$$
$$H_2 = \{\{1,3\},\{2,4\},\{5,7\},\{6,8\}\}$$

$K_i(E) = \{\omega | h_i(\omega) \subseteq E\}$, knowledge of event E of player i: all states of the world where i knows event E (whether E or not E)

$$K_1(\{2,6\}) = \{2,6\},\ K_1(\{2,4\}) = \emptyset,$$
$$K_1(\{2\}) = \emptyset,\ K_1(\{2,4,6\}) = \{2,6\},$$
$$K_1(\{2,4,6,8\} = \{2,4,6,8\}$$

$K_I(E) = \{\omega | \cup h_i(\omega) \subseteq E\}$: mutual knowledge of E

$$K_{1,2}(\{2,6\}) = \emptyset, K_{1,2}(\{2,4,6,8\}) = \{2,4,6,8\}$$

$K_I^2(E) = \{\omega | \cup h_i(\omega) \subseteq K_I(E)\}$: mutual knowledge of mutual knowledge of E

$$K_{1,2}^2(\{2,4,6,8\}) = \{2,4,6,8\}$$

$$K_1(\{2,4,6\}) = \{2,6\}, K_2(\{2,4,6\}) = \{2,4\},$$
$$K_{1,2}(\{2,4,6\}) = \{2\},$$
$$K_{1,2}^2(\{2,4,6\}) = K_{1,2}(K_{1,2}(\{2,4,6\})) = K_{1,2}(\{2\}) = \emptyset,$$
$$K_1(K_{1,2}(\{2,4,6\})) = K_1(\{2\}) = \emptyset,$$

if true state is 2, $\{2,4,6\}$ is mutual knowledge,
but neither 1 nor 2 know that it is mutual knowledge!

Not hat example: $\hat{H}_2$: $\{1,3\},\{2,4\},\{5,7,\mathbf{8}\},\{6\}$

$$K_1(\{2,4,6\}) = \{2,6\}, K_2(\{2,4,6\}) = \{2,4,6\},$$
$$K_{1,2}(\{2,4,6\}) = \{2,6\},$$
$$K_{1,2}^2(\{2,4,6\}) = K_{1,2}(\{2,6\}) = \{6\},\ K_{\{1,2\}}^3 = K_{\{1,2\}}(\{6\}) = \emptyset$$

$K_I^\infty(E)$: common knowledge of E

$$K_{1,2}^\infty(\{2,4,6,8\}) = \{2,4,6,8\}$$
$\{2,4,6,8\}$ is common knowledge

Definition: H is a *meet* of H_i and H_j, $H = H_i \wedge H_j$: smallest coarsening (Partition H for which H_i and H_j are refinements and any refinement of which is not refinement of either H_i or H_j)

$$\hat{H} = H_1 \wedge H_2 = \{\{1,3,5,7\},\{2,4,6,8\}\}$$
$$H_3 = \{\{1,2\},\{3,4\},\{5,6\},\{7,8\}\}$$

Aumann's Agreeing to Disagree theorem

Theorem: Suppose that it is common knowledge at ω that player i's posterior probability of event E is q_i and that player j's posterior probability of E is q_j. Then $q_i = q_j$.

Let $H = H_i \wedge H_j$, or the meet of H_i and H_j

The posterior probability for player i of event E given the partition element $\omega \in h_i^k$ is given by:

$$p(E|h_i^k) = \frac{P(E \cap h_i^k)}{p(h_i^k)}$$

Let us now sum over all partition elements h_i^k that are inside the meet element h, doing so for both players. This yields:

$$\sum_k p(E|h_i^k) = \frac{\sum P(E \cap h_i^k)}{\sum p(h_i^k)} = \frac{P(E \cap h)}{p(h)}$$
$$\sum_k p(E|h_j^k) = \frac{\sum P(E \cap h_j^k)}{\sum p(h_j^k)} = \frac{P(E \cap h)}{p(h)}$$

Notice that, if priors $p(h)$ are equal, posteriors will also be equal for all players. This occurs because of the assumption that the posteriors are common knowledge. If this is true, then each player must have already incorporated all knowledge that the other player has into their own posterior probability. Given the same priors, posterior beliefs must then be the same.

No Trade Theorem

If two people cannot disagree about the probabilities, they could not have a purely speculative trade.

Theorem: Fix some random variable X and a state $\omega \in \Omega$. It cannot be common knowledge at ω that both players have strictly positive conditional expectations from the bet.

How one could explain speculative trades?

1. "Liquidity'' traders (or other sorts of irrational behaviour). Importantly, a (rational) trader may be unsure whether one is trading with rational trader or liquidity trader, hence common knowledge will not result.

2. Assume away common priors. If there is no common prior, a trader may infer that "his own part is unfavorable'', but may also infer that another trader has a different prior.

Topic 2: Rationality

Defining Rationality

$$s_i \text{ is rational if } \exists\sigma_{-i}\ \forall\sigma_i' \in S_i, u(s_i, \sigma_{-i}) \geq u(\sigma_i', \sigma_{-i})$$

- The mixed strategy σ_{-i} is said to rationalize the play of s_i
- Note that σ_{-i} should technically be interpreted not as the other player's strategy, but as your belief about their strategy
- Rationality by itself does not restrict beliefs about other players: rational player could hold any belief about opponent's play σ_{-i}. But, given that belief, she should best respond

S is a rational strategy if there exists a single mixed strategy σ_{-i} of the other player for any possible alternative strategy S', such that the utility from playing S if the other player plays σ_{-i} is greater than (or equal to) the utility from playing any other strategy S'. This definition just means that a strategy is rational if it may be useful – if it might be good to play it against some particular mixed strategy of the other play. Many of my strategies could be rational if the other player has many possible mixed strategies.

Summary: A strategy is rational if there exists some (i.e. at least one) mixed strategy of the other player for which it is the best response.

Strictly Dominated Strategies

A pure strategy s_i is strictly dominated for player i if there exists σ_i' such that:

$$u_i(\sigma_i', s_{-i}) > u_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}$$

This means that a strategy is strictly dominated if there is some other mixed strategy that would always get you a higher payoff for any possible strategy that other guy could play.

Undominated Strategies

$$s_i \text{ is undominated if } \forall\sigma_i' \in \Sigma_i\ \exists\sigma_{-i}: u(s_i, \sigma_{-i}) \geq u(\sigma_i', \sigma_{-i})$$

Strategy S is undominated if for every other possible strategy s' there exists some mixed strategy of the other player σ_{-i}, such that the utility from playing S if the other player plays σ_{-i} is greater than (or equal to) the utility from playing the other strategy σ'.

Summary: A strategy is undominated if every single other mixed strategy I might play yields a worse outcome for me in at least one potential circumstance (i.e. for at least one particular strategy of the other player).

Two Player Equivalence Theorem

Theorem: In a game with two players, a strategy s_i is rational if and only if it is undominated.

Proof: the 'only if' part is pretty clear – a dominated strategy is never going to be rational. To prove the 'if' (i.e. all undominated strategies are rational in at least some possible circumstance), we use the supporting hyperplane theorem.

Hyperplane Definition

A hyperplane is a concept in geometry. It is a generalization of the plane into a different number of dimensions. A hyperplane divides a space into two half-spaces. A hyperplane is said to support a set S if:

- S is entirely contained in one of the two closed half-spaces determined by the hyperplane
- S has at least one point on the hyperplane.

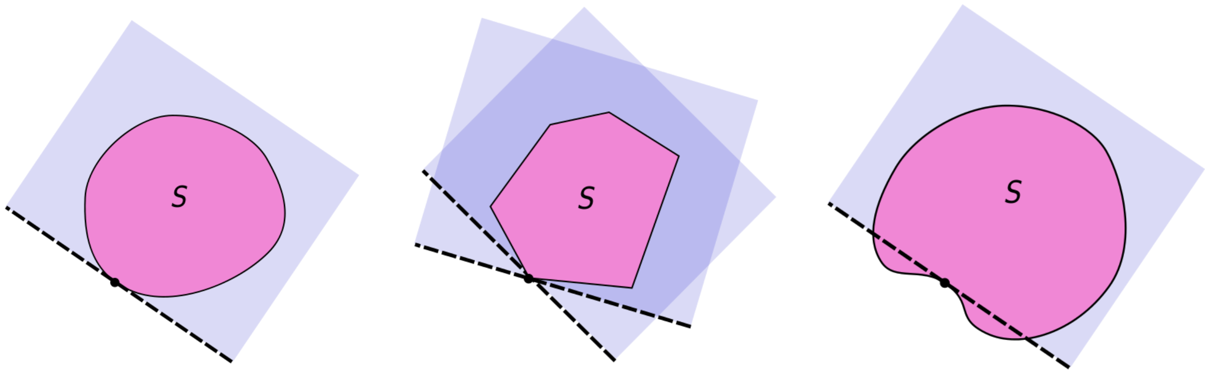

Supporting Hyperplane Theorem

If a set S is closed and convex, then there exists a supporting hyperplane at every boundary point of

- A set is closed if and only if it contains all of its limit points (i.e. boundary points)
- An object is convex if for every pair of points within the object, every point on the straight line segment that joins them is also within the object

Application

Undominated strategies correspond to points on the boundary of S, as these are points that do not always have some other point that is 'further out' than themselves. A strategy is rational if a hyperplane at that point can support S, which would imply that this point is the best strategy to play given some particular mixed strategy of the other player (i.e. given some particular hyperplane). The supporting hyperplane theorem states that any point on the boundary of a closed, convex set can be a supporting hyperplane, and thus any undominated strategy (boundary point) is also rational (on a hyperplane). Note the diagram on the far right, where a boundary point cannot host a supporting hyperplane – this is because the set S is not convex.

Rationalizability

A strategy is rationalizable if a perfectly rational agent could justifiably play it against perfectly rational opponents. More precisely, assume that a player plays an expected payoff maximizing action given some belief about the other player's action, which must be an optimal action given some belief about the other players' action and so on. If an action can be rationalized by such an infinite sequence of reasoning, we say that the action is rationalizable.

A rationalizable strategy profile is a strategy profile that consists only of rationalizable strategies. Nash equilibrium strategies are always rationalizable. In two-player games, rationalizable strategies are simply those that survive the iterated elimination of strictly dominated strategies. In n-agent games, this isn't so. Rather, rationalizable strategies are those that survive iterative removal of strategies that are never a best response to any strategy profile by the other agents. If we change the definition of rationalizability to include correlation, proposition 2 is correct for N > 2, as this permits the strategy space S to again become convex.

Note that rationalizable strategies are a subset of iteratively undominated strategies, so we can try to find rationalizable strategies through iterative elimination of dominated strategies. If the structure of the game and the rationality of the players are common knowledge, then both players will choose an iteratively undominated strategy.

Solution Concept 1: Iterated Elimination of Dominated Strategies

Once we have eliminated dominated strategies for each player, it often turns out that a pure strategy that was not dominated at the outset is now dominated. Thus, we can undertake a second round of eliminating dominated strategies. Indeed, this can be repeated until pure strategies are no longer eliminated in this manner. In a finite game, this will occur after a finite number of rounds and will always leave at least one pure strategy remaining for each player. If strongly (resp. weakly) dominated strategies are eliminated, we call this the iterated elimination of strongly (resp. weakly) dominated strategies.

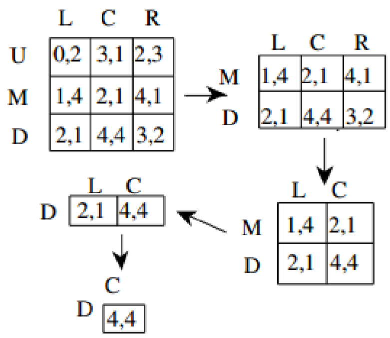

Figure 4.1. The iterated elimination of strongly dominated strategies

1.5 Back to the Cournot Example

We now apply the technique of iterated elimination of strictly dominated strategies to the Cournot Competition. (Our terminology in the following example is somewhat loose and informal.) In the first step, we note that both firms must choose a quantity between $[0, \infty]$. We denote this:

$$A_1^1 = [0, \infty]$$
$$A_2^1 = [0, \infty]$$

Notice, however, that it is not rational for player 1 to choose any quantity that is outside of the range $[0, 1/2]$ since player 1's best response function is 0 outside of that range. Therefore, playing *any* strategy which is defined over $[0, \infty]$ is dominated by playing one over $[0, 1/2]$. The same reasoning holds for player 2. Thus, at the second iteration, we can argue that the pair of best responses must be

$$A_1^2 = [0, 1/2]$$
$$A_2^2 = [0, 1/2]$$

Given that player 2 only plays in the range $[0, 1/2]$, then player 1 can *restrict* his best response function to only these values. In Figure 3, this is depicted with the dotted lines. Consider the point where the horizontal orange line intersects $B_1(s_2)$. Since player 2 will only play strategies below the dashed line, then player 1 need only consider strategies between $[1/4, 1/2]$. The same situation holds for player 2.

$$A_1^3 = [1/4, 1/2]$$
$$A_2^3 = [1/4, 1/2]$$

By iterating the best response function, a formal argument can be made to show that the limit of this process will converge to the point which equals the sum:

$$\frac{1}{2} - \frac{1}{4} + \frac{1}{8} - \frac{1}{16} \cdots = \frac{1}{3}$$

Topic 3: Nash Equilibrium

Solution Concept 2: Pure Strategy Nash Equilibrium

Strategies s^* are a (pure strategy) Nash equilibrium if for each player i and for any strategy $s_i \in S_i$:

$$u(s_i^*, s_{-1}^*) \geq u(s_i, s_{-1}^*)$$

Suppose that each player is rational and knows his own payoff function, and that the strategy choice of the players are mutually known. Then these choices constitute a Nash equilibrium in the game being played. This version does not require any assumption of common knowledge of rationality, but instead requires that all players make correct conjectures about other players' actions.

One can alternatively regard this version of the Nash equilibrium as the result of deductive consideration by both players, a stable equilibrium outcome over many plays, or as a self-enforcing agreement between players.

"If players think and share the belief that there is an obvious way to play a game, then it must be a Nash equilibrium."

Solution Concept 3: Mixed Strategy Nash Equilibrium

Strategies σ^* are a mixed strategy Nash equilibrium if for each player i and for any strategy $\sigma_i \in S_i$:

$$u(\sigma_i^*, \sigma_{-1}^*) \geq u(\sigma_i, \sigma_{-1}^*)$$

Note that utility in this case is simply expected utility from playing the given mixed strategy, with all the possible payoffs weighted by their probability and then summed. Hence this can also be written:

$$\sum u_i(s_i^*, s_{-i})\sigma_{-1}(s_{-i}) \geq \sum u_i(s_i, s_{-i})\sigma_{-1}(s_{-i})$$

Mixed strategy equilibria can be found by selecting the probability distribution over pure strategies that makes the opponent indifferent between their possible strategies, and then doing this for every player. By ensuring that the other player is indifferent in this way, you ensure that they do not know what you are going to do.

Two-Person Game

Suppose that the structure of the game, the rationality of the players, and their conjectures are mutual knowledge. Then the conjectures constitute a mixed-strategy Nash equilibrium.

N-Person Game

Suppose that the structure of the game, the rationality of the players, and their conjectures are mutual knowledge. Suppose also that for each player, all other players hold the same conjecture about that player. Then the conjectures constitute a mixed-strategy Nash equilibrium.

Interpretations of Mixed Strategies

- Randomize to confuse your opponent
- Randomize when you are uncertain about the other's action
- A description of what might happen in repeated play: count of pure strategies in the limit
- Strategies are not random, but based upon some deterministic process the other player cant observe
- Mixed strategies describe population dynamics

Example 1

Suppose player 1 plays U with probability π

	L	R
U	5,1	0,0
D	4,4	1,5

$u_1(1, \pi) = 5\pi$

$u_1(0, \pi) = 4\pi + 1 \cdot (1 - \pi)$

If proper mixed strategy, $u_1(1, \pi) = u_1(0, \pi) \Rightarrow 5\pi = 4\pi + 1 - \pi \Rightarrow \pi = \frac{1}{2}$

NE: $(0.5, 0.5)$

Note: Player 1 is indifferent between U and D and choose probabilities **to make the other player indifferent**

Then just do the same thing for the other player.

Example 2

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Let player 2 play B with p, F with $1 - p$.
- If player 1 best-responds with a mixed strategy, player 2 must make him indifferent between F and B (why?)

$$u_1(B) = u_1(F)$$
$$2p + 0(1 - p) = 0p + 1(1 - p)$$
$$p = \frac{1}{3}$$

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Likewise, player 1 must randomize to make player 2 indifferent.
  - Why is player 1 willing to randomize?
- Let player 1 play B with q, F with $1 - q$.

$$u_2(B) = u_2(F)$$
$$q + 0(1 - q) = 0q + 2(1 - q)$$
$$q = \frac{2}{3}$$

Existence of Nash Equilibrium

Defining Correspondences

A correspondence is a mapping of units of one set to units of another set. It can have multiple inputs and multiple outputs. A best response correspondence gives the optimal action for a player as a function of the strategies of all other players. If there is always a unique best action given what the other players are doing, then this is a function. If for some opponent's strategy, there is a set of best responses that are equally good, then this is a correspondence.

A mixed strategy profile σ^* is a Nash equilibrium if and only if it is a fixed point of the best response correspondence, i.e. $\sigma^* \in \sigma_i(\sigma_{-i})$.

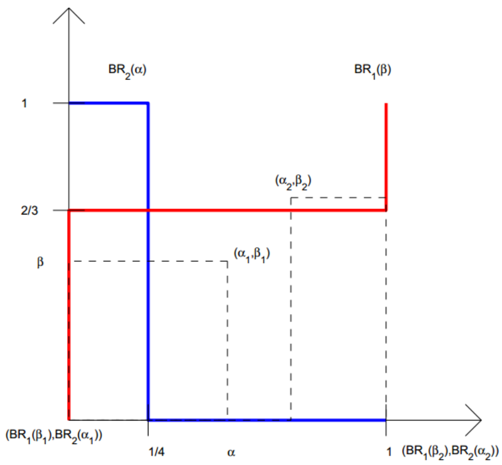

Fixed Point Theorem

Kakutani's fixed point theorem guarantees that a correspondence $\sigma_i(\sigma_{-i}): \Sigma \to \Sigma$ has a fixed point if the following four conditions are satisfied:

- Σ (the strategy set of all players) is compact (closed and bounded, containing its boundary), convex, and nonempty (satisfied in all finite games)
- $\sigma_i(\sigma_{-i})$ is nonempty (satisfied in any actual game)
- $\sigma_i(\sigma_{-i})$ is convex (satisfied for any mixed strategy)
- $\sigma_i(\sigma_{-i})$ is upper hemicontinuous (will hold if u_i is continuous and compact)

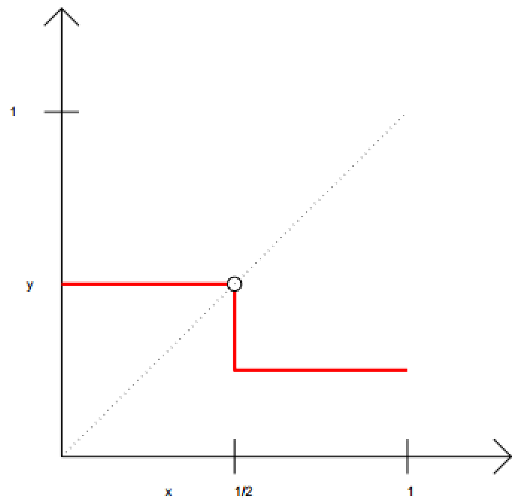

; $r(\frac{1}{2})$ is the convex set $[\frac{1}{4}, \frac{1}{2})$ but that this set is not closed.

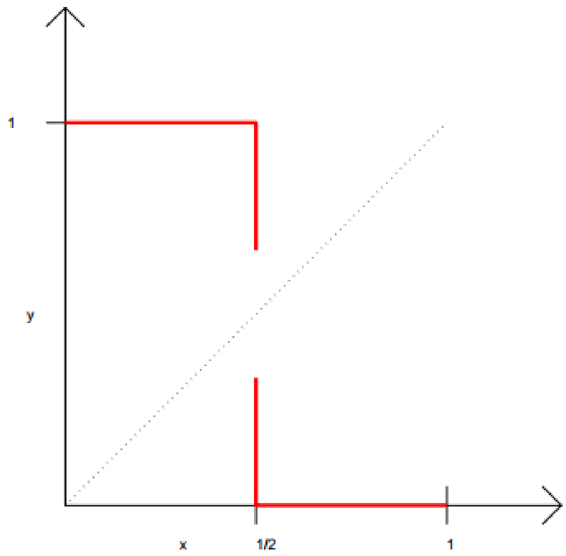

Theorem

All finite games (finite number of players with a finite number of strategies) have at least a mixed strategy Nash equilibrium, because in these games the BRF will always have a fixed point.

Advantages and Disadvantages of Nash Equilibria

Advantages

- Mixed strategy Nash equilibrium always exists
- NE is a long-used and well understood tool

Disadvantages

- Pure strategies: need to know opponent's strategies
- Mixed strategies 2 person: need to know opponent's conjectures (randomisations)
- Mixed strategies N person: opponent's conjectures must be common Knowledge
- NE sensitive to small changes in payoffs
- Nash equilibria are not necessarily self-enforcing agreements
- NE need not be unique

	Current	Modern	New
Current	10,10	10,5	0,9
Modern	5,10	5, 5	8,5
New	9, 0	5, 8	7,7

Solution Concept 4: Correlated Equilibria

By using a publically observable random variable, players can receive any payoff vector within the convex-hull of Nash equilibrium payoffs. Note that every Nash equilibrium is also a correlated equilibria, which means that correlated equilibria always exist in finite games. Mixed strategy Nash equilibria are a subset (special case) of correlated equilibria that do not permit randomization of actions based upon external observed signals.

$$\sum u_i(s_i^*, s_{-i})\sigma_{-1}(s_{-i}|s_i) \geq \sum u_i(s_i, s_{-i})\sigma_{-1}(s_{-i}|s_i)$$

We see here that the definition of a correlated equilibria is the same as for a mixed strategy Nash equilibria, except that probabilities can now be conditioned upon the action of the other player (this can also be interpreted as being conditional upon some external signal). In the case of a mixed strategy Nash equilibrium, strategies must be independent, and so this conditionality is forbidden.

A game of chicken:

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

NE: $(U, R), (D, L), ((\frac{2}{3}, \frac{1}{3}), (\frac{2}{3}, \frac{1}{3}))$

Let probabilities of states be given by $(\alpha + \beta + \gamma + \delta = 1)$:

α	β
γ	δ

For a correlated equilibrium to occur, each player's actions U, D must coincide with the opponent's expectations of their actions U_{-i}^e, D_{-i}^e. Note that payoffs come from the strategy actually played, while the probabilities come from the strategy your opponent expects you to play.

$$u_1(U_1|U_2^e) = 6\alpha + 2\beta$$
$$u_1(D_1|U_2^e) = 7\alpha + 0\beta$$
$$\therefore 6\alpha + 2\beta \geq 7\alpha$$
$$2\beta \geq \alpha$$

$$u_1(D_1|D_2^e) = 7\gamma + 0\delta$$
$$u_1(U_1|D_2^e) = 6\gamma + 2\delta$$
$$\therefore 7\gamma \geq 6\gamma + 2\delta$$
$$\gamma \geq 2\delta$$

$$u_2(L_1|L_1^e) = 6\alpha + 2\gamma$$
$$u_2(R_1|L_1^e) = 7\alpha + 0\gamma$$
$$\therefore 6\alpha + 2\gamma \geq 7\alpha$$
$$2\gamma \geq \alpha$$

$$u_2(L_1|R_1^e) = 7\beta + 0\delta$$
$$u_2(R_1|R_1^e) = 6\beta + 2\delta$$
$$\therefore 7\beta \geq 6\beta + 2\delta$$
$$\beta \geq 2\delta$$

This set of inequalities defines the set of all correlated equilibria for this game. If the inequalities are inconsistent (i.e. cannot all simultaneously hold), then there are no correlated equilibria. Example:

$$\alpha + 2\left(\frac{\alpha}{2}\right) + \frac{\alpha}{4} = 1$$
$$\alpha\left(2 + \frac{1}{4}\right) = 1$$
$$\alpha = \frac{4}{9} \therefore \beta = \gamma = \frac{2}{9}, \delta = \frac{1}{9}$$

Solution Concept 5: Bayesian Nash Equilibrium

A strategic-form Bayesian game consists of:

- A set of players N
- A set of states Ω
- And for each player $i \in N$:
    - A set of available actions A_i
    - A set of signals T_i and a signal function that associates a signal with each state $(\tau_i: \Omega \to T_i)$
    - A prior belief about the states (a probability measure p_i on Ω)
    - Preferences $\succcurlyeq_i$ over outcomes

A strategy profile $(\sigma_1^*, \dots, \sigma_n^*)$ is a Bayesian Nash equilibrium if, for every player i, for every type t_i of i and for any strategy $\sigma_i(t_i)$:

$$E[u_i(\sigma_i^*, \sigma_{-1}^* | t_i)] \geq E[u_i(\sigma_i, \sigma_{-1}^* | t_i)]$$

In brief, in a Nash equilibrium of a Bayesian game each player chooses the best action available to him given the signal that he receives and his belief about the state and the other players' actions.

Topic 4: Extensive Form Games

Introduction

An extensive form game consists of:

- a set of players
- a tree
- an allocation of each non-terminal node to a player
- an information partition
- a payoff at each terminal node

An information set is a collection of nodes such that he same player moves at each node in the collection, and the same moves are available at each node in the information set, and these nodes cannot be distinguished by the player to which they belong. When each information set is a singleton, this is a perfect information game.

Strategies in an extensive form game must assign a probability to every action available at every information set, even those not on the equilibrium path.

Extensive form games map uniquely into strategic (matrix) form, while strategic form games do not. This may be an indication that extensive form games preserve irrelevant information.

A pure strategy Nash equilibrium will always exist in a finite, perfect information extensive form game.

Solution Concept 6: Subgame Perfect Nash Equilibrium

A subgame of an extensive form game Γ is an extensive form game constructed from Γ which:

- Begins at a node that is in a singleton information set, which is not a terminal node
- Includes all nodes and only nodes that follow the beginning node
- Preserves the tree structure and information partition of the original game
- Note: A subgame which is not a game itself is proper subgame.

A Nash equilibrium is subgame perfect if the player's strategies constitute a Nash equilibrium in every subgame.

When you cannot draw a full game tree to solve for SPNE, you have to structure a SPNE argument on smaller pieces of the tree, doing each piece-by-piece and then putting these together for an overall solution. One way to do this is by using the one deviation principle.

Problems with SPNE

Subgame perfect Nash Equilibria require players to assume that other players are perfectly rational even if they have already played non-rational strategies (i.e. off the equilibrium path). It seems unreasonable to require that players maintain this belief despite evidence to the contrary, but SPNE provides no way of dealing with this. Another criticism of SPNE is that it dramatically fails to predict behaviour in situations like the centipede game, where people tend to cooperate until near the end of the game, rather than defect immediately as SPNE would suggest.

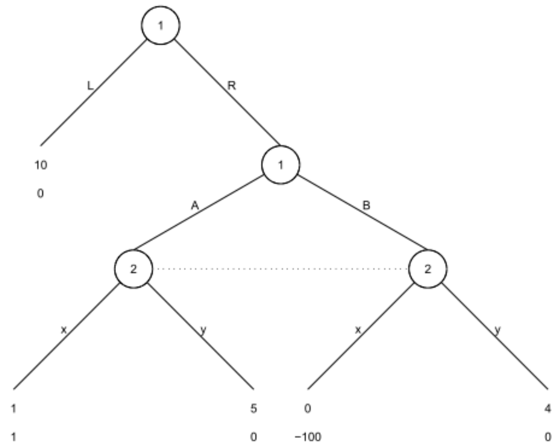

Here (L, A, x) is the unique SPE. However, player 2 has to put a lot of trust into player 1's rationality in order to play x. He must believe that player 1 is smart enough to figure out that A is a dominant strategy in the subgame following R. However, player 2 might have serious doubts about player 1's marbles after the guy has just foregone 5 utils by not playing L.[1]

One Deviation Principle

In a finite game, strategy profile s is subgame perfect if and only if it satisfies the one deviation principle, which states that for each player i and each information set of i:

- Fix strategies (not actions) of other players s_{-i}
- Fix moves of i at all other information sets
- If i then cannot improve her conditional payoff at this information set by deviating from s_i at this information set only, this strategy satistifes the one deviation principle

To deal with infinite games, we can simply solve for the discount factor that would equate payoffs under the s_i and s_i', thereby seeing if/when s_i' is a profitable deviation.

Example 1

	C	D
C	4, 4	0, 5
D	5, 0	1, 1

TfT strategy says:

Play C in the first period. In period $t > 1$ play whatever the other player did in period $t - 1$.

history that ends with the outcome (C, D). If they both stick to TfT we will have the path

$$...(C, D) \mid (D, C), (C, D), (D, C), (C, D), (D, C)....$$

The continuation payoff for the player 1 here is

$$5 + 0 \cdot \delta + 5\delta^2 + 0 \cdot \delta^3 + 5\delta^4 + ... = \frac{5}{1 - \delta^2}$$

But what if player 1 does single deviation after the history ending with (C, D)? This means that instead of repeating D of player 2 he plays C and then continues doing whatever TfT was prescribing. We fix player 2's strategy and get the following path:

$$...(C, D) \mid (C, C), (C, C), (C, C), (C, C), (C, C)....$$

This gives player 1 the payoff

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

Let's see for which δ it is profitable to deviate:

$$\begin{aligned} \frac{4}{1-\delta} &> \frac{5}{1-\delta^2} \quad \Rightarrow \\ \delta &> 1/4 \end{aligned}$$

For any δ bigger than $\frac{1}{4}$ the guy will prefer to deviate. So, TfT against TfT is not SPNE for high δ.

Example 2

	C	D
C	(4,4)	(0,5)
D	(5,0)	(1,1)

Payoff: "average payoff": $(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \pi_t$

(Trigger) strategy: $\sigma_i^t =$ start with C; play C if (C,C) and play D otherwsie (if (D,C),(C,D) or (D,D) are ever observed).

Deviation to D in period 1: (note why player 1 plays D,D,D)

$$U(D, D, D \ldots, \sigma_2) = 5 + \delta \times 1 + \delta^2 \times 1 + \ldots \qquad = 5 + \frac{\delta}{1 - \delta}$$

$$U(\sigma) = 4 + 4\delta + 4\delta^2 + \ldots \qquad = 4 + 4\frac{\delta}{1 - \delta}$$

Not profitable for large $\delta < 1$.

After D: (D, D, D, ...), (D, D, D, ...)

$$U(C, D, D \ldots, \sigma_2) = 0 + \delta \times 1 + \delta^2 \times 1 + \ldots \qquad = 0 + \frac{\delta}{1 - \delta}$$

$$U(D, D, D, \sigma) = 1 + 1\delta + 1\delta^2 + \ldots \qquad = 1 + \frac{\delta}{1 - \delta}$$

Not profitable for any $\delta < 1$ (mention (D, C, C))

Backward Induction

This is effectively a way of performing iterated elimination of weakly dominated strategies, but using a particular order. It is an algorithm to find subgame perfect Nash equilibria in a perfect information game.

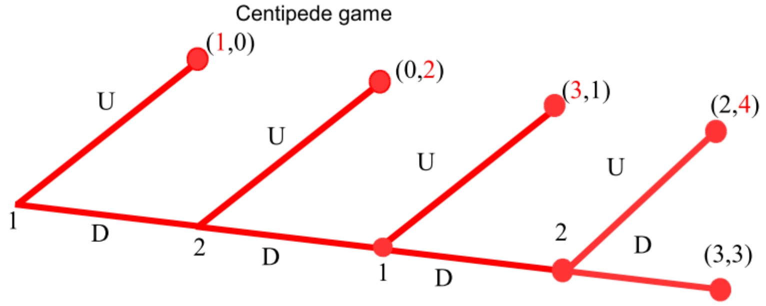

	U,U	U,D	D,U	D,D
U,U	1,0	1,0	1,0	1,0
U,D	1,0	1,0	1,0	1,0
D,U	0,2	0,2	3,1	3,1
D,D	0,2	0,2	2,4	3,3

Repeated Games

An infinitely repeated game is defined by the following elements. The stage game is played at each discrete time period t = 1,2,..., and at the end of each period the action choice of each player is revealed to everybody. A history in time period t is simply a sequence of action profiles from period 1 through period $t-1$, e.g.:

$$\big((C,C),(C,C),(C,D)\big)$$

The Folk Theorems

Friedman Theorem

The main property of infinitely repeated games is that the set of equilibria becomes very large as players get more patient $(\delta \to 1)$. Given any payoff vector that gives each player more than some Nash equilibrium outcome of the stage game, for sufficiently large values of δ, there exists some subgame perfect equilibrium that yields the payoff vector at hand as the average value of the payoff stream. This is stated formally below.

Let G be a finite, static game of complete information. Let e_i denote the payoffs from a Nash equilibrium of G for player i, and let x be a feasible payoff of G (a convex combination of pure strategy payoffs). For all $x_i > e_i$, there exists δ^* such that for every $\delta^* < \delta < 1$, there exists a subgame perfect equilibrium of $G(\infty, \delta)$ that achieves x_i as the average payoff for player i.

Example

	C	D
C	(4,4)	(0,5)
D	(5,0)	(1,1)

NE is (D,D) giving payoffs $e=(1,1)$. Suppose we want to achieve payoff:

$$x_i = \frac{1}{2}u_i(C,C) + \frac{1}{2}u_i(C,D) = \frac{4}{2} + \frac{5}{2} = \frac{9}{2} > e_i$$

Consider the following strategies:

σ_1: start with C, play C if player 2 played C in all odd periods and D in all even periods, play D otherwise

σ_2: start with C, play C in all odd periods and D in all even periods if player 1 played C in all periods, else play D

These strategies constitute a SPNE

Classic Folk Theorem
The difference between this theorem and the Friedman theorem is that it applies to all equilibria, not just those with payoffs higher than the Nash equilibria in the stage game.

Let r_i be player i's reservation utility, or minmax value. This is defined as:

$$r_i = \min_{-i}\left[\max_{i} u_i(i, -i)\right]$$

In other words, the minmax value is the smallest payoff that the other player can force you to accept.

Let G be a finite static game of complete information. For any feasible payoff x such that for all i, $x_i > r_i$, there exists $\bar{\delta}$ such that for every $\delta < \delta < 1$, there exists a Nash equilibrium of $G(\infty, \delta)$ that achieve x as the average payoff.

To be achievable, payoffs must be both feasible (coloured in), and individually rational (blue).

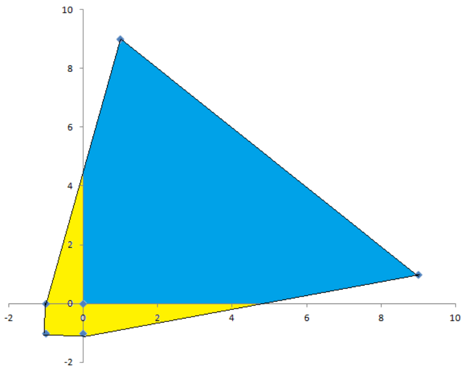

Example

	C	D	N
C	(4,4)	(0,5)	(0,0)
D	(5,0)	(1,1)	(1,0)
N	(0,-2)	(0,-2)	(0,-2)

Suppose we want to support the payoffs $(5{,}0)$:
σ_1: start with D and play D if player 2 plays C, play N otherwise
σ_2: start with C and play C every turn
These strategies constitute a Nash equilibrium, though not a subgame perfect one

Fudenberg and Maskin Theorem
Let G be a finite static game of complete information. For any feasible payoff x such that $x > r$, and if the set of feasible payoffs has as many dimensions as there are players, then there exists $\bar{\delta}$ such that for every $\delta < \delta < 1$, there exists a subgame perfect Nash equilibrium of $G(\infty, \delta)$ that achieve x as the average payoff.

Example 1

	C	D	N
C	(4,4)	(0,5)	(0,0)
D	(5,0)	(1,1)	(1,0)
N	(0,-2)	(0,-2)	(0,-2)

Player 2 has incentives to deviate to D. If 2 deviates, player 1 plays N for one period and then D again (not profitable for player 2 to deviate). But this is not a credible punishment, so we need to reward player 1 for punishing. We can do this if player 2 plays C rather than D for one period (note: player 2 payoff is the same). Summing up: on the equilibrium they play (D, N). If player 2 deviates, they play (N, N) (or (N, D)), then (D, C) and then continue with (D, N). If player 1 deviates from N, then play (D, N).

Example 2

	C	D	N
C	(4,4)	(0,0)	(-1,-1)
D	(0,0)	(1,1)	(-1, -1)
N	(-1,-1)	(-1,-1)	(-2,-2)

Note that in this example the theorem does not hold because the payoff space has only one dimension. This means that it is not possible to reward/punish one player without also rewarding/punishing the other.

Strategies as Automata

An automata is an abstraction of how players play and make decisions in infinite repeated games.

An automaton for player i in game $G(\infty, \delta)$ with actions (x_i, x_{-i}) includes:

- A set of states Ω
- An initial state $\omega_0 \in \Omega$
- An output function $s(\omega_i)$ what assigns a set of actions for each state
- A transition function $t(\omega_i, (x_i, x_{-i}))$ that assigns a state to every combination of state and actions

Example 1

Example 141.1 (*A machine for the "grim" strategy*) The machine $\langle Q_i, q_i^0, f_i, \tau_i \rangle$ defined as follows is the simplest one that carries out the ("grim") strategy that chooses C so long as both players have chosen C in every period in the past, and otherwise chooses D.

- $Q_i = \{\mathcal{C}, \mathcal{D}\}$.
- $q_i^0 = \mathcal{C}$.
- $f_i(\mathcal{C}) = C$ and $f_i(\mathcal{D}) = D$.
- $\tau_i(\mathcal{C}, (C, C)) = \mathcal{C}$ and $\tau_i(\mathcal{X}, (Y, Z)) = \mathcal{D}$ if $(\mathcal{X}, (Y, Z)) \neq (\mathcal{C}, (C, C))$.

Example 2

	C	D
C	(4,4)	(0,5)
D	(5,0)	(1,1)

Consider the SPNE strategies written as:

σ_1: start with C, play C if player 2 played C in all odd periods and D in all even periods, play D otherwise

σ_2: start with C, play C in all odd periods and D in all even periods if player 1 played C in all periods, else play D

Written as automata these strategies become:

- $\Omega = \{\Sigma_1, \Sigma_2, \Delta\}$
- $\omega_0 = \Sigma_1$
- $s(\Sigma_1) = (C, C), s(\Sigma_2) = (C, D), s(\Delta) = (D, D)$
- $t(\Sigma_1, (C, C)) = \Sigma_2, t(\Sigma_1, (x_i, x_{-i}) \neq (C, C)) = \Delta, (\Sigma_2, (C, D)) = \Sigma_1$

Solution Concept 7: Weak Perfect Bayesian Equilibrium

Motivation and Problems with SPNE

Although SPNE is good at ruling out non-credible threats, it is not very helpful in ruling out some other equilibria that we think are unreasonable. For example, the strategy $\sigma_1 = \text{Out}, \sigma_2 = (F, F)$ is subgame perfect Nash equilibrium, however it is also ridiculous. We want a better solution concept that can rule out these sort of equilibria.

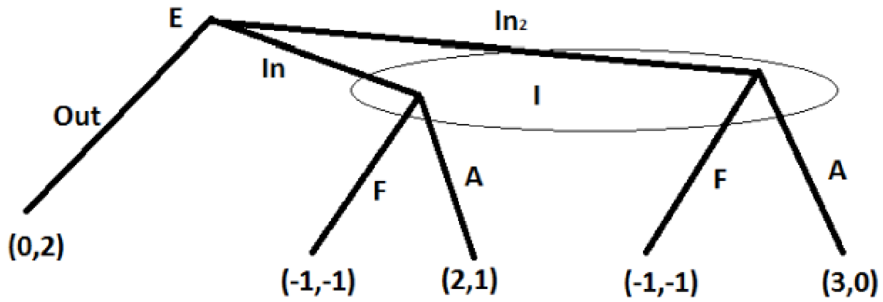

Sequential Rationality

A system of beliefs μ in extensive form game G is a specification of probability $\mu(x) \in [0, 1]$ for each decision node x in G such that for each information set H, $\sum_{x \in H} \mu(x) = 1$. Let i be a player who moves at an information set H. A strategy profile $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ in extensive form game G is sequentially rational at information set H given a system of beliefs μ if for all σ_i^*:

$$E(u_i|H, \mu, \sigma_i, \sigma_{-i}) \geq E(u_i|H, \mu, \sigma_i^* \sigma_{-i})$$

If a strategy profile σ_i satisfies this condition for all information sets H, then we say that σ is sequentially rational given belief system μ_i.

Weak Perfect Bayesian Equilibrium

A profile of strategies and the system of beliefs (σ, μ) is a weak perfect Bayesian equilibrium in extensive form game G if it has the following properties:

- The strategy profile σ is sequentially rational given belief system μ
- Beliefs are *Bayesian consistent* with this strategy. This means that for any information set H_i of player i such that $P(H_i|\sigma) > 0$, and for all nodes $x \in H_i$, it must be the case that $\mu(x) = \frac{P(x|\sigma)}{P(H_i|\sigma)}$
- Note that the sum of probabilities across the nodes within any information set must be one

The $P(x|\sigma)$ term is very important. It means if strategy profile σ implies that one player always plays a particular action x, then $P(x|\sigma) = 1$. Values less than one are possible in the event of mixed strategies, or if there is some probabilistic action by nature. Note that a weak PBE need not be subgame perfect.

On Finding Solutions

There is no standard algorithm for finding PBE. Some heuristics:

- start by eliminating all the dominated and iteratively dominated strategies for each player
- every time you learn something about the informed player's strategy, see what it implies for the uninformed player's beliefs
- every time you revise the uninformed player's beliefs, see what it implies for her strategy
- every time you learn something about the uninformed player's strategy, see what it implies for the informed player's strategy

Example 1

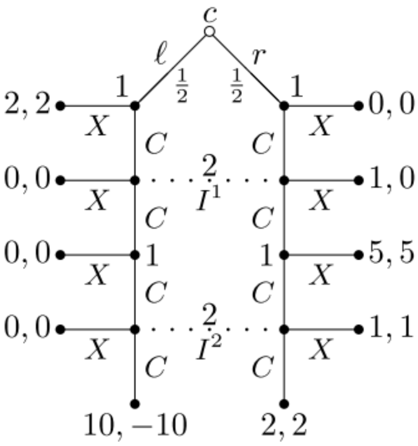

EXAMPLE 236.1 In any sequential equilibrium of the game in Figure 236.1

- player 1 chooses C after the history r
- player 1 chooses X after the history (r, C, C)
- player 2 chooses C at his information set I^1
- player 2 chooses X with probability at least $\frac{4}{5}$ at his information set I^2 (otherwise player 1 chooses C after the histories ℓ and (ℓ, C, C), so that player 2 assigns probability 1 to the history (ℓ, C, C, C) at his information set I^2, making C inferior to X)
- player 1 chooses X after the history ℓ.

Thus player 2's belief at I^1 assigns probability 1 to the history r while his belief at I^2 assigns positive probability to chance having chosen ℓ (otherwise C is better than X).

Example 2

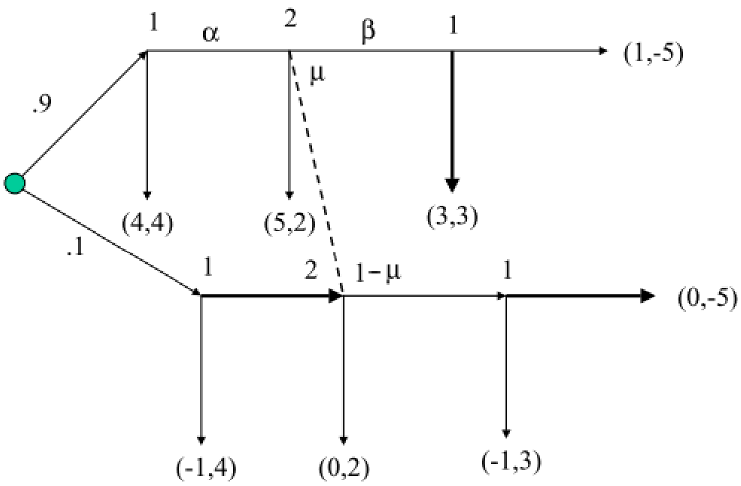

- Sequential rationality requires that at the last note in the upper branch player 1 goes down, and at the last node of the lower branch player 1 goes across. Moreover, it requires that player 1 goes across at the first node of the lower branch
- Suppose that player 1 goes down with probability 1 at the first node on the upper branch. Then, by Bayes' rule, player 2 must assign probability 1 to the lower branch at his information set (as you cant get to that information set on the top branch given this strategy profile), and hence he must go down with probability 1. In that case, it is better for player 1 to go across and get 5, rather than going down and getting 4 –a contradiction.
- Now, suppose that player 1 goes across with probability 1 at the first upper node. Then by Bayes' rule, player 2 must assign probability .9 to the upper branch in his information set. Hence, if he goes down, he gets 2; if he goes across, he gets .9 × 3 + .1 × (− 5) = 2.2. Then, he must go across with probability 1. In that case, player 1 must go down with probability 1 at the first upper node – another contradiction
- Therefore, player 1 must mix at the first upper node. In order to have this, player 1 must be indifferent between going across and going down. Player 2 mixes so as to make player 1 indifferent:
- For indifference: $4 = 5 \times (1-\beta) + 3 \times \beta \Rightarrow 4 = 5 - 2\beta \Rightarrow \beta = \frac{1}{2}$
- Player 1 must also mix so as to make player 2 indifferent:
- For indifference: $0.9 \times \alpha \times 3 + 0.1 \times 1 \times -5 = 0.9 \times \alpha \times 2 + 0.1 \times 1 \times 2 \Rightarrow 0.9\alpha = 0.7 \Rightarrow \alpha = \frac{7}{9}$
- For Bayesian consistency to hold: $\mu = \frac{0.9\alpha}{0.9\alpha + 0.1(1)} = \frac{0.9(7/9)}{0.9(7/9)+0.1} = \frac{6.3}{7.2} = \frac{7}{8}$

Example 3

Although we have introduced the concept of weak PBE here in the context of dynamic games of incomplete information the definition applies to any game and particularly to any game of imperfect information (after all the Harsanyi trick translates incomplete information into imperfect information). Let us apply the concept to the following game:

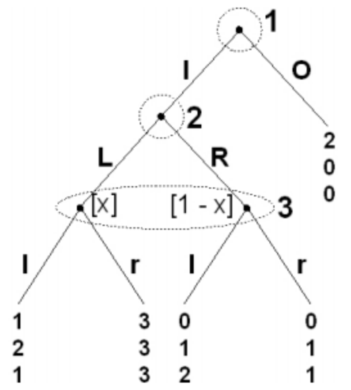

	l	r
L	2, 1	3, 3
R	1, 2	1, 1

- this game has one subgame other than the game itself
- it begins at player *2*'s singleton information set and corresponds to the simultaneous move game shown below (player *2* is the row player)
- the unique NE in this subgame is (L, r), so the unique SPNE of the entire game is $(s_1^*, s_2^*, s_3^*) = = (I, L, r)$
- this strategy profile together with the belief $x = 1$ is also a weak PBE of this game
- however, there is another weak PBE where $(s_1^*, s_2^*, s_3^*) = (O, L, l)$ and $x = 0$
- thus, a **weak PBE need not be subgame perfect**
- the problem is that player 3's belief is inconsistent with player 2's strategy (and that **weak PBE does not impose any restrictions on off-the-equilibrium-path beliefs**).

Example 4

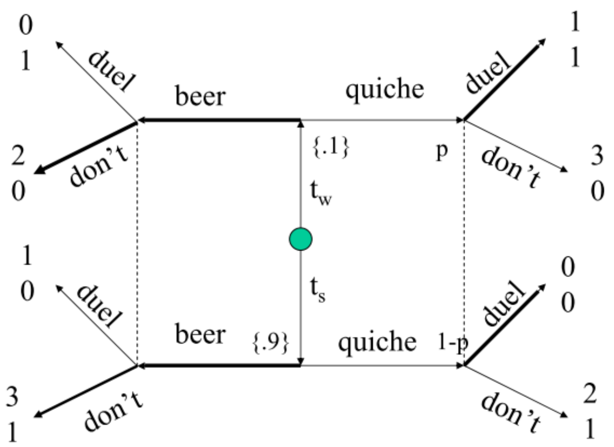

p=μ(t_w|quiche) > 1/2, and
p=1/2 & Prob (duel | quiche) ≥1/2

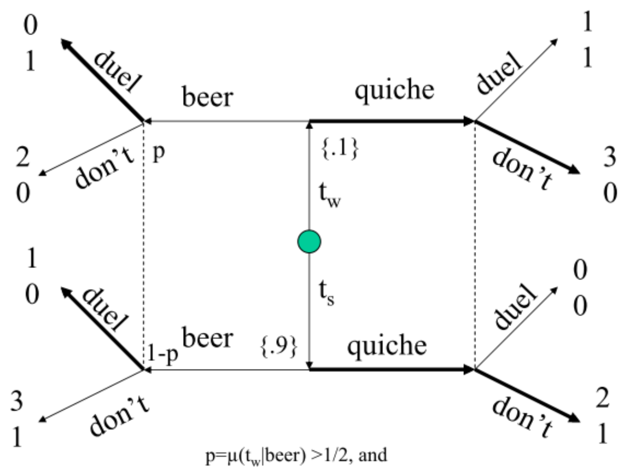

p=μ(t_w|beer) >1/2, and
p=1/2 & Prob (duel | beer) ≥1/2

A Mixed SE

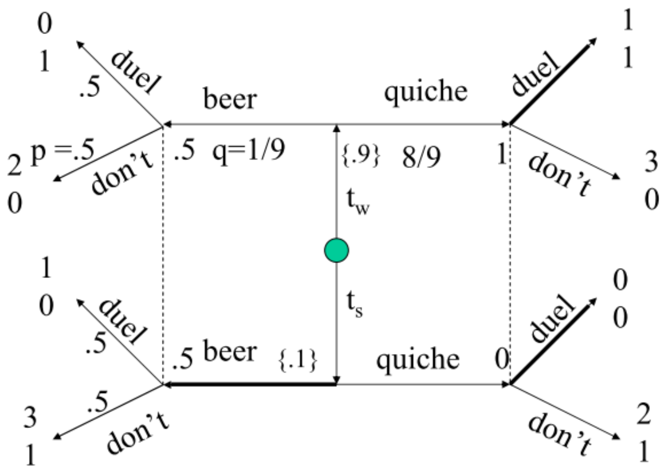

Solution Concept 8: Sequential Equilibrium

This was developed to remedy the deficiency of WPBE that off-equilibrium beliefs can be irrational. It is a much stronger concept than WPBE.

A strategy profile and system of beliefs (σ, μ) is a sequential equilibrium of extensive form game G if:

- The strategy profile σ is sequentially rational given belief system μ
- Beliefs are *consistent* with this strategy. There exists a sequence of completely mixed strategies (all actions played with positive probability) with $\lim_{k\to\infty}(\sigma^k) = \sigma$ such that $\lim_{k\to\infty}(\mu^k) = \mu$, where μ^k are the beliefs derived from strategy profile σ^k using Bayes' rule

In every sequential equilibrium (σ, μ) of an extensive form game G, the equilibrium strategy profile σ constitutes a subgame perfect equilibrium of G. Every finite extensive game with perfect recall has a sequential equilibrium.

Example 1

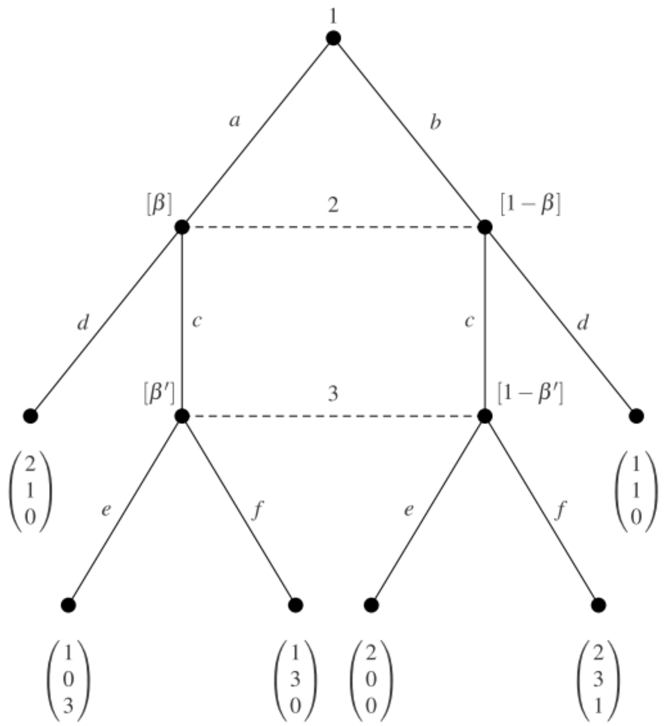

Consider a single case where $\beta > \frac{1}{4}$

- Observe that consistency requires the beliefs of players 2 and 3 to be the same, so $\beta = \beta'$
- Given that $\beta > \frac{1}{4}$, $E(e) \geq 0.4 \times 3 + 0.6 \times 1 \geq 1.8$, $E(f) \leq 0.4 \times 0 + 0.6 \times 1 \leq 0.6$, $\therefore p_3(e) = 1$
- For player 2, $E(d) = 0.4 \times 1 + 0.6 \times 1 = 1$, $E(c) = 0.4 \times 0 + 0.6 \times 0 = 0$, $\therefore p_2(d) = 1$
- Clearly, player 1 will play a, hence we can determine that the true value of β must be 1
- Thus, this sequential equilibria is defined by: $p_1(a) = 1, p_2(c) = 0, p_3(e) = 1, \beta = \beta' = 1$

Consider a second case where $\beta = \frac{1}{4}$

- For this to be a consistent belief, it must hold that $p_1(a) = \frac{1}{4}$
- For this to be true, player 1 must be indifferent between a and b, meaning that:

$$2 \times p_2(d) + p_2(c)\big(1 \times p_3(e) + 1 \times p_3(f)\big) = 1 \times p_2(d) + p_2(c)\big(2 \times p_3(e) + 2 \times p_3(f)\big)$$

$$2p_2(d) + (1 - p_2(d))(1) = 1p_2(d) + (1 - p_2(d))(2)$$
$$p_2(d) + 1 = -p_2(d) + 2$$
$$p_2(d) = \frac{1}{2}$$

- Given that player 2 is now also mixing, they too must be indifferent between d and e:

$$1 \times \frac{1}{4} + 1 \times \frac{3}{4} = \frac{1}{4} \times (0 \times p_3(e) + 3 \times p_3(f)) + \frac{3}{4} \times (0 \times p_3(e) + 3 \times p_3(f))$$
$$1 = 3 \times p_3(f)$$
$$p_3(f) = \frac{1}{3}$$

- Thus, this sequential equilibria is defined by: $p_1(a) = \frac{1}{4}, p_2(c) = \frac{1}{2}, p_3(e) = \frac{2}{3}, \beta = \beta' = \frac{1}{4}$

Example 2

Definition (Consistency): An assessment (μ, β) is **consistent** if there exists a completely mixed sequence (μ^n, β^n) that converges to (μ, β) such that μ^n is derived from β^n using Bayes' rule for all n.

An assessment (μ, β) is a **sequential equilibrium** if it is sequentially rational and consistent. To illustrate, consider the game in Figure 11.3 again. Let μ be the probability assigned to the node that follows L, and consider the assessment $((A, L, L'), \mu = 0)$. For this to be a sequential equilibrium, we have to find a completely mixed behavioral strategy profile β^n such that

$$\beta_1^n(A) \to 1, \beta_2^n(L) \to 1, \beta_3^n(L') \to 1, \mu^n = \frac{\beta_2^n(L)}{\beta_2^n(L) + \beta_2^n(R)} \to 0,$$

which is not possible. However, the assessment given by $((D, L, R'), \mu = 1)$ is easily checked to satisfy sequential rationality. To check consistency, let

$$\beta_1^n(D) = 1 - \frac{1}{n}, \beta_2^n(L) = 1 - \frac{1}{n}, \beta_3^n(R') = 1 - \frac{1}{n}, \mu^n = 1 - \frac{1}{n}.$$

Notice that μ^n is derived from β^n via Bayes' rule and $(\mu^n, \beta^n) \to (\mu, \beta)$. Therefore, this assessment is a sequential equilibrium.

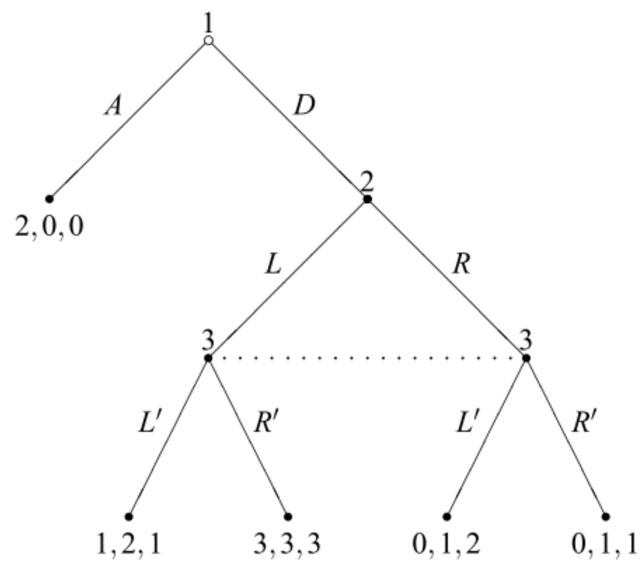

Figure 11.3: PBE may have "unreasonable" beliefs

The Intuitive Criterion

The Intuitive Criterion is a special technique for equilibrium refinement in signaling games. The basic idea is that off-equilibrium beliefs should be reasonable in some sense. Specifically, a WPBE fails the intuitive criteria if it places non-zero probability on a message m_t that is equilibrium dominated for type t.

Given a PBE in a signalling game, the message m_t of player type t is equilibrium dominated if t's equilibrium payoff $U_t^* > \max U_t(m, a)$. That is, a message is equilibrium dominated if deviating from that equilibrium and sending an alternate message would always yield a strictly lower payoff.

If the information set is off the equilibrium path and m is equilibrium-dominated for type t, then the receiver's belief $\mu(t|m)$ should place probability zero on type t (as they would never want to deviate and send this message).

"I'm having beer, and you should infer from this that I'm surly, for so long as it is common knowledge that you will not fight if I eat quiche, I would have no incentives to drink beer and make this speech if I were wimpy."

The only pooling SE that survives
the Intuitive Criterion