GAME THEORY

(based largely on Steven Matthews' notes)

1. Concepts

1.a. Basic definitions

Game theory is the study of strategic interaction – that is, interactions where a player's payoffs depend on others' actions as well as own her own. Game theory can also be thought of as multi-person decision theory, and many of the ideas from our study of choice under uncertainty will show up.

As an introduction, we'll consider two-player games.

<u>Example:</u>

Player's 1's utility payoffs (NOT wealth payoffs) are given in the following matrix:

		Rain God: *Sun*	Rain God: *Rain*
Player 1	*No Umbrella*	5	0
	Umbrella	1	3

In general, we have two expected-utility maximizing players.

Player 1's set of possible actions $A = \{a_1, \dots , a_K\}$,

Player 2's set of possible actions $B = \{b_1, \dots , b_J\}$.

Player 1's utility is given by $u(a, b)$, where u: $A \times B \rightarrow \boldsymbol{R}$. (We'll add Player 2's payoffs later.)

1.b. Dominance and belief rationality

Dominance:

Action a "(strictly) dominates" action a' if

$u(a, b) > u(a', b)$ for all $b \in B$.

Action a "weakly dominates" action a' if

$u(a, b) \geq u(a', b)$ for all $b \in B$, and

$u(a, b) > u(a', b)$ for some $b \in B$.

Action a is "undominated" (or "admissible") if it is not weakly dominated by any $a' \in A$.

Action a is "dominant" if it (strictly) dominates every other action $a' \in A$.

Action a is "weakly dominant" if it weakly dominates every other action $a' \in A$.

A rational player will never choose a strictly dominated action. She will always pick a dominant action if there is one.

Examples:

		Player 2	
		C	*D*
Player 1	*C*	–1	–4
	D	0	–3

		Player 2	
		L	*R*
	T	3	0
Player 1	*M*	1	1
	B	0	3

Beliefs:

Player 1's beliefs about Player 2's actions are given by a probability distribution $p = [p(b_1), \dots, p(b_J)]$, where each $p(b_j) \geq 0$ and $\sum_{j=1}^{J} p(b_j) = 1$. The set of such beliefs is $\Delta(B)$.

Thus, the expected utility $u(a, p)$ of an action a is given by $u(a, p) = \sum_{j=1}^{J} p(b_j) u(a, b_j)$.

(Note that the expected utility function $u(a, p)$ maps from $A \times \Delta(B) \rightarrow \boldsymbol{R}$, while our original utility function $u(a, b)$ maps from $A \times B \rightarrow \boldsymbol{R}$.)

A rational player will always choose an action that maximizes expected utility given beliefs p. Since the set of actions A is finite, there is always at least one best action.

Action a is a "(weak) best response" to beliefs p if

$$u(a, p) \geq u(a', p) \text{ for every action } a' \in A.$$

The "best response correspondence" $BR(p)$ is the set of best responses:

$$BR(p) = \arg\max_{a \in A} u(a, p).$$

Example: Sun-Rain game.

Let p = prob(Sun). Then

$u(No\ Umbrella, p) = 5p + 0(1 - p) = 5p$, and

$u(Umbrella, p) = 1p + 3(1 - p) = 3 - 2p$.

Thus, $BR(p) = \begin{cases} \{NoUmbrella\} & p > 3/7 \\ \{NoUmbrella, Umbrella\} & p = 3/7 \\ \{Umbrella\} & p < 3/7. \end{cases}$

Relating dominance and belief rationality:

Action a is "never a weak best response" (NWBR) if it is not a best response to any belief $p \in \Delta(B)$.

A rational player will not choose a NWBR.

Proposition 1.b.1: A strictly dominated action is a NWBR.

Proof: If action a dominates action a', then by definition $u(a, b) > u(a', b)$ for all $b \in B$, and so $u(a, p) > u(a', p)$ for all $p \in \Delta(B)$. Thus, a' is a NWBR. *Q.E.D.*

The converse of the proposition is not true. Here is a counterexample:

Player 2

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

Let p = prob(L). Then

$$BR(p) = \begin{cases} \{T\} & p > 1/2 \\ \{T, B\} & p = 1/2 \\ \{B\} & p < 1/2. \end{cases}$$

So action M is a NWBR, but M is undominated.

However, consider a "mixed action," where Player 1 flips a coin to decide between T and B. Call that action F.

The expected utility $u(F, p) = 0.5\, u(T, p) + 0.5\, u(B, p)$

$= 0.5[3p] + 0.5[3(1 - p)] = 0.5[3] = 1.5.$

Since $u(F, p) = 1.5$ for all p, in particular $u(F, L) = u(F, R) = 1.5$.

Thus, F dominates M.

1.c. Mixed actions

A "mixed action" is a probability distribution $\sigma \in \Delta(A)$:

$$\sigma = [\sigma(a_1), \ldots, \sigma(a_K)].$$

The "support" of a mixed action is the set of pure actions given positive probability:

$$\text{Supp}(\sigma) = \{a \in A : \sigma(a) > 0\}.$$

It will sometimes be convenient to think of pure actions as just a special case of mixed actions.

The expected utility $u(\sigma, p)$of a mixed action σ given beliefs p is

$$u(\sigma, p) = \sum_{a \in A} \sigma(a) u(a, p) \left(= \sum_{a \in A} \sum_{b \in B} \sigma(a) p(b) u(a, b) \right).$$

(Note that the expected utility function for a mixed action $u(\sigma, p)$ maps from $\Delta(A) \times \Delta(B)$ $\to$ $\boldsymbol{R}$.)

There is a sense in which mixed actions "have no value." A mixed action cannot yield a payoff higher than the best pure action in its support, since the payoff to the mixture is a convex combination of the payoffs to the pure actions.

We can redefine the best responses as mixed actions:

$$BR(p) = \arg\max_{\sigma \in \Delta(A)} u(\sigma, p).$$

Proposition 1.c.1: A mixed action σ is a best response if and only if every pure action in its support is a best response.

Proof:

$$u(\sigma, p) = \sum_{a \in A} \sigma(a) u(a, p) = \sum_{a \in \text{Supp}(\sigma)} \sigma(a) u(a, p) \leq \max_{a \in \text{Supp}(\sigma)} u(a, p). \qquad \textit{Q.E.D.}$$

Example: Sun-Rain game.

Let p = prob(*Sun*), $\sigma = [\sigma(\textit{No Umbrella}), \sigma(\textit{Umbrella})]$. Then

$$BR(p) = \begin{cases} \{(1,0)\} & p > 3/7 \\ \Delta(A) & p = 3/7 \\ \{(0,1)\} & p < 3/7. \end{cases}$$

Note that $BR(3/7)$ includes (1, 0) and (0, 1).

Dominance by a mixed action:

Action $a \in A$ is "(strictly) dominated by a mixed action" $\sigma \in \Delta(A)$ if

$u(a, b) < u(\sigma, b)$ for all $b \in B$.

Proposition 1.c.2: An action $a \in A$ is dominated (possibly by a mixed action) if and only if it is a NWBR.

Proof: The proof of the "only if" direction is almost identical to the proof of Proposition 1.b.1. To prove the "if" direction, we show that if action a is not strictly dominated, then it is a best response for some belief.

Let $x^a = (u(a, b_1), \dots, u(a, b_J)) \in \mathbf{R}^J$. Define sets X and Y as follows:

$X = \{x \in \mathbf{R}^J \mid x = (u(\sigma, b_1), \dots, u(\sigma, b_J))$ for some $\sigma \in \Delta(A)\}$, and

$Y = \{y \in \mathbf{R}^J \mid y >> x^a\}$.

Note that $x^a \in X$, that X and Y are disjoint (since a is not strictly dominated), and that x^a lies on the boundary of Y. Both X and Y are convex, and Y is open. A separating hyperplane theorem ensures that there is a nonzero vector $p \in \mathbf{R}^J$ such that $(y - x^a) \cdot p > 0$ for all $y \in Y$, and $(x - x^a) \cdot p \le 0$ for all $x \in X$. Since the set Y is unbounded above, p must be nonnegative. We can thus interpret p as a probability vector. (Just normalize by dividing through by $\sum p_j$.) Since $(x - x^a) \cdot p \le 0$ for all $x \in X$, a is a weak best response to belief p: any other action gives a weakly lower expected payoff. *Q.E.D.*

2. Normal (Strategic) Form Games

2.a. Definitions

Definition: A "normal form game" G is a collection of three things: players, actions for each player, and utility functions for each player.

There are N players.
The "action set" (or "strategy set") for Player i is denoted A_i.
The set of "action profiles" is $A = A_1 \times A_2 \times \ldots \times A_N$.
Player i's utility function is $u_i : A \to \mathbf{R}$.

<u>Example:</u> Sun-Rain game.

We'll add the Rain God's payoffs. The Rain God likes the sun and also likes looking at umbrellas.

		Rain God	
		Sun	*Rain*
Player 1	*No Umbrella*	5, 3	0, 0
	Umbrella	1, 5	3, 2

REMEMBER: Payoffs are not outcomes. They give the *utility* from outcomes. They already include risk aversion, concern for others, etc. So a rational player will maximize his or her own expected utility.

2.b. Solution concepts

How will people play a given game? We will examine different ideas of what a solution, or "equilibrium," should be.

2.b.i Equilibrium in undominated strategies

We already know that a rational player will not play a dominated action. Sometimes that fact by itself is enough to give us a prediction.

<u>Example:</u>

		Player 2: *C*	Player 2: *D*
Player 1	*C*	−1, −1	−4, 0
	D	0, −4	−3, −3

This game is the Prisoners' Dilemma.
Each player has only one action that is not strictly dominated: *D*. So (*D*, *D*) is the solution.

<u>Example:</u> Second-price auction.

Bidding your valuation is a weakly dominant strategy. If you bid higher, then (ignoring ties) either you don't change the outcome, or you win and pay a price greater than your valuation. If you bid lower, then (again ignoring ties) either you don't change the outcome, or you lose when you could have won at a price below your valuation.

Looking for a solution in undominated strategies (especially in strategies that are not strictly dominated) is a very reasonable thing to do, but often players have multiple undominated strategies. In the Sun-Rain game, for example, both of Player 1's actions are undominated.

2.b.ii. Iterated deletion of strictly dominated actions

Example: Sun-Rain game.

Player 1 has no dominated actions. However, for the Rain God *Sun* is dominant. If we remove *Rain* from the game, then *No Umbrella* becomes dominant for Player 1. So if both players are rational, and if Player 1 *knows* that the Rain God is rational, and that Player 1 knows the Rain God's payoffs, then we predict that (*No Umbrella*, *Sun*) will be played.

The order of deletion does not matter. That is, if more than one action is strictly dominated, it doesn't matter whether they are deleted simultaneously or sequentially, or in what order.

This solution concept requires not only that players are rational, but that 1) they know each other's payoffs, and they know that they know each other's payoffs, and they know that they know that they know each other's payoffs, … , and 2) they know that the other is rational, and they know that they know that the other is rational, … , for as many levels as there are rounds of deletion.

Example:

		Player 2	
		L	*R*
Player 1	*U*	8, 10	–1000, 9
	D	7, 6	6, 5

An event is "common knowledge" if both players know that it is true, and both know that both know that it is true, and both know that both know that both know that it is true, and so on infinitely.

Again, many games are not solvable this way. That is, in many games more than one action profile survives iterated deletion of strictly dominated strategies.

Alternatively, we could try iteratively deleting *weakly* dominated strategies.

Example:

		Player 2: *S*	Player 2: *D*
Player 1	*T*	1, 1	2, 1
	B	0, 1	1, −1000

The order of deletion does matter.

There are still many games that are not solvable this way.

This solution concept is less reasonable than iterated deletion of strictly dominated actions, because a rational player might choose a weakly dominated action. In the example above, *D* was ruled out only because *S* is better than *D* when Player 1 chooses action *B*, which is strictly dominated. But shouldn't Player 2 "know" that Player 1 will never play a dominated strategy?

As another alternative, we could try iteratively deleting NWBR actions. For two-player games, that approach yields exactly the same solution as the iterated deletion of strictly dominated actions. (The strategy profiles that survive are called "rationalizable.") With more than two players, things are more complicated. In defining a NWBR, should we allow for *correlated* randomization by the other players, or should we require that the others' strategies are independent?

2.b.iii. Nash equilibrium

A pure action profile $a^* = (a_1^*, \ldots, a_N^*)$ is a "Nash equilibrium" if for each player i,

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \text{ for every } a_i \in A_i.$$

That is, $a_i^* \in BR(a_{-i}^*)$ for every player i.

A (mixed) action profile $\sigma^* = (\sigma_1^*, \ldots, \sigma_N^*)$ is a "Nash equilibrium" if for each player i, $\sigma_i^* \in BR(\sigma_{-i}^*)$. That is, $a_i \in BR(\sigma_{-i}^*)$ for every $a_i \in \text{Supp}(\sigma_i^*)$.

<u>Example:</u> The Prisoners' Dilemma.

		Player 2	
		C	D
Player 1	C	−1, −1	−4, 0
	D	0, −4	−3, −3

The unique Nash equilibrium is (D, D). That profile is also the only one to survive iterated deletion of strictly dominated actions.

Proposition 2.b.1: If σ^* is a Nash equilibrium, then σ^* survives iterated deletion of strictly dominated actions. (That is, every action profile $a \in \text{Supp}(\sigma^*)$ survives.)
Proof: Practice.

2.c. Finding equilibria

2.c.i. Calculating mixed strategy equilibria

Example: "Battle of the Sexes"

		Woman	
		Fight	*Ballet*
Man	*Fight*	2, 1	0, 0
	Ballet	0, 0	1, 2

There are two pure strategy Nash equilibria: (*Fight*, *Fight*) and (*Ballet*, *Ballet*).
We can also find an equilibrium in mixed strategies:
If Man is mixing, then he must be indifferent between playing *Fight* and playing *Ballet*, which implies conditions on the strategy that *Woman* is playing.
Let $p = \sigma_{Woman}(Fight)$. Then

$U_{Man}(Fight) = 2p + 0(1 - p) = 2p$, and

$U_{Man}(Ballet) = 0p + 1(1 - p) = 1 - p$.

Indifference requires that $2p = 1 - p$, so $p = 1/3$.
Similarly, Woman must be indifferent between *Fight* and *Ballet* to be willing to mix, which implies (by a similar calculation) that $\sigma_{Man}(Fight) = 2/3$.
So each player is willing to mix only if the other player is mixing. That result implies that there is no Nash equilibrium where one player mixes and the other plays a pure strategy. The unique mixed strategy equilibrium is thus ((2/3, 1/3), (1/3, 2/3)).

In general, games may have mixed strategy equilibria where some players mix and some don't, or where a player mixes over some of her actions and not others. In looking for mixed strategy equilibria, it is important to check all the cases. For large games, there are a lot of cases.[1]

[1] Life is full of sadness.

2.c.ii. Existence of Nash equilibria

Definition: A correspondence f mapping from $A \subseteq \mathbf{R}^N$ to a closed set $Y \subseteq \mathbf{R}^M$ is "upper hemicontinuous" if

1) the set $\{(x, y) : x \in A, y \in f(x)\}$ is closed (with respect to $A \times Y$), and
2) for any compact set $A^C \subseteq A$, the set $\{y : y \in f(x) \text{ for some } x \in A^C\}$ is compact.

Part 1) says that the correspondence has "closed graph." Part 2) says that the image of a compact set is compact.

Claim: The best response correspondence is upper hemicontinuous.
Suppose that $\{p_n\} \to p$ is a sequence of mixed strategies by the other players. If an action $a \in BR(p_n)$ for all n, then $a \in BR(p)$.

Proposition 2.c.1: Every finite game (that is, every game with finite sets of players and actions) has at least one (possibly mixed) Nash equilibrium.
Proof idea: We apply Kakutani's fixed point theorem to the best response correspondence. A Nash equilibrium is an action profile that is a best response to itself – that is, a fixed point of the best response correspondence. $BR(\cdot)$ is a nonempty, convex-valued (why?), upper hemicontinuous correspondence mapping from a nonempty, convex, compact set (the set of mixed action profiles) to itself, so Kakutani's fixed point theorem guarantees the existence of a fixed point.

We have seen that the Nash equilibrium may not be unique.

Nash equilibrium has *nothing* to do with Pareto optimality.

The Nash equilibrium correspondence (mapping from payoffs to actions) is also upper hemicontinuous.

2.d. Examples

2.d.i. Classic games

<u>Example:</u> "Matching Pennies"

		Player 2	
		Heads	*Tails*
Player 1	*Heads*	1, –1	–1, 1
	Tails	–1, 1	1, –1

Matching pennies is a "zero sum game." It has no Nash equilibrium in pure strategies. There is a unique mixed strategy equilibrium: ((½, ½), (½, ½)).

<u>Example:</u> "Divide the Dollar"

There are two players, with action sets $A_1 = A_2 = \mathbf{R}_+$. Payoffs are as follows:

$$u_i(a_i, a_j) = \begin{cases} a_i, & \text{if } a_i + a_j \leq 1 \\ 0, & \text{otherwise.} \end{cases}$$

Divide the Dollar is not a finite game, so Proposition 2.c.1 does not guarantee that a Nash equilibrium exists. In fact, though, there are many. Any pair (a_1, a_2) such that either $a_1 + a_2 = 1$ or $\min\{a_1, a_2\} \geq 1$ is a pure strategy equilibrium. Note that some of those equilibria involve both players choosing weakly dominated actions. There are also many mixed strategy equilibria.

<u>Example:</u> Cournot duopoly

There are two firms, each of whom chooses a quantity $q_i \in A_i = \mathbf{R}_+$. Define Q as the sum of q_1 and q_2. Payoffs are as follows:

$$u_i(q_i, q_j) = q_i P(Q) - c_i(q_i),$$

where $P(\cdot)$ is inverse demand and $c_i(\cdot)$ is firm i's cost function.
For the sake of a simple example, let $P(Q) = a - Q$, and let $c_i(q_i) = c\, q_i$ for each i, where $a > c$. Then the best response correspondences are given by

$$BR_i(q_j) = \text{argmax } q_i\,(a - q_i - q_j) - c\, q_i.$$

At an interior solution, the first order condition is $a - 2q_i - q_j - c = 0$, so

$$q_i(q_j) = \frac{a - c - q_j}{2} \text{ if } \frac{a - c - q_j}{2} \geq 0, \text{ and 0 otherwise. Similarly,}$$

$$q_j(q_i) = \frac{a - c - q_i}{2} \text{ if } \frac{a - c - q_i}{2} \geq 0, \text{ and 0 otherwise.}$$

The pure strategy Nash equilibrium, then, is $q_i^* = q_j^* = \frac{a-c}{3}$. Total quantity $Q^* = \frac{2(a-c)}{3}$, and the equilibrium price $P^* = \frac{1}{3}a + \frac{2}{3}c$.

With N firms, each $q_i^* = \frac{a-c}{N+1}$. Total quantity $Q^* = \frac{N(a-c)}{N+1}$, and the equilibrium price $P^* = \frac{1}{N+1}a + \frac{N}{N+1}c$. As N grows, those values approach the competitive price c and quantity $(a - c)$.

2.d.ii. Thought-provoking examples

Suppose that a game has two Nash equilibria. One might think that if one equilibrium involves weakly dominated strategies and the other doesn't, then the first is less likely to be played. One might also think that an equilibrium that is Pareto dominated by another equilibrium is unlikely to be played. But consider the following example:

Player 2

Player 1		*L*	*R*
	T	10, 0	5, 2
	B	10, 11	2, 11

That game has two pure strategy equilibria: (*B*, *L*) and (*T*, *R*). The first equilibrium involves both players choosing weakly dominated actions, but it strictly Pareto dominates the second!

Or consider the following three-player game, where Player 1 chooses rows, Player 2 chooses columns, and Player 3 chooses the box:

Box 1

	X	*Y*
A	1, 1, 0	1, 0, 1
B	1, 1, 1	0, 1, 1

Box 2

	X	*Y*
A	1, 0, 1	1, 1, 0
B	1, 1, 0	0, 1, 0

The only pure strategy Nash equilibrium, (*B*, *X*, *Box* 1), has Player 1 playing a weakly dominated action (*B*).

3. Extensive Form Games

Remember the Battle of the Sexes:

		Woman: *Fight*	Woman: *Ballet*
Man	*Fight*	2, 1	0, 0
	Ballet	0, 0	1, 2

What if Man, working downtown, is farther from the boxing arena and from the ballet theater, so that he has to leave work and make his decision first?

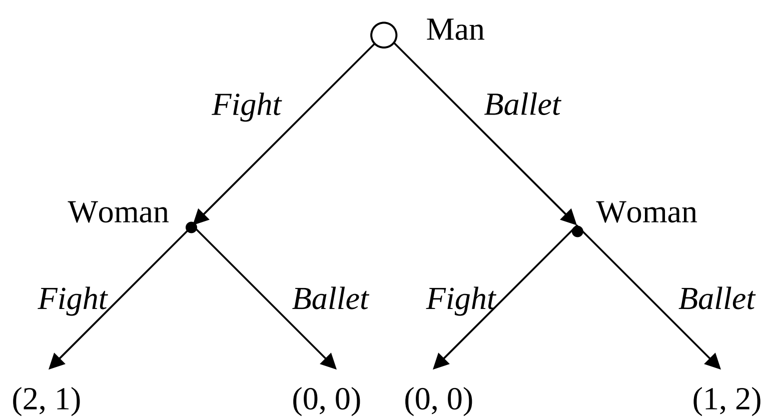

But that picture does not quite capture the situation that we're interested in. We need a way to indicate who knows what when, as well as who moves when.

3.a. Definitions

3.a.i. Extensive-form game ingredients

1) Players 0, 1, … , N. Player 0 is Nature.
2) Nodes: $y \in Y$.
   a. Decision nodes $x \in X$.
   b. Terminal nodes $z \in Z$.
3) Directions: predecessors and successors.
   a. Each node (except the *initial node* y_0) has one *immediate predecessor*.
   b. Terminal nodes have no *successors*, immediate or otherwise.
4) The *game tree* is initial node y_0 and all its successors.
   a. A *subtree* is any node and all its successors.
5) Whose decision: a map $i: X \to N$ that specifies which player gets to move at each decision node.
6) Actions: $A(x)$ is the set of actions available to player $i(x)$ at decision node x.
   a. Each different action leads to a different immediate successor node.
7) Information sets: H is a partition of decision nodes into *information sets*. A player cannot tell which node within an information set she's at, although she can distinguish between information sets.
   a. $h(x)$ is the information set that contains decision node x.
   b. The same player moves at each decision node in an information set: if $x' \in h(x)$, then $i(x) = i(x')$. So we can write $i(h)$.
   c. The same actions are available at each decision node in an information set: if $x' \in h(x)$, then $A(x) = A(x')$. So we can write $A(h)$.
8) Nature's moves: Nature moves randomly according to commonly known probabilities.
9) Payoffs: A set of maps $u_i: Z \to \mathbf{R}$ that specify a player's utility at each terminal node.

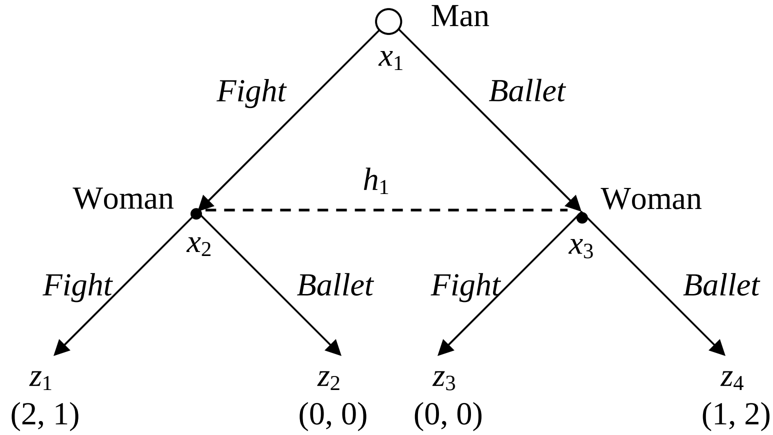

3.a.ii. Properties

1) A game has *perfect recall* if players don't forget anything (e.g., their own previous actions) during the game. (We could define this property formally using restrictions on information sets. But we won't.)
2) A game has *perfect information* if each information set contains only one node – that is, if players know the history of moves so far.
3) A game has *complete information* if the structure of the game is common knowledge – that is, each player could draw the game tree. (E.g., players know each others' payoffs.)
    a. We'll usually transform incomplete information games into imperfect information games by letting Player 0 choose the structure randomly and (at least to some players) unobservably.

3.a.iii. Strategies and payoffs

In extensive form games, we will distinguish between strategies and actions. Actions were defined above.

Let H_i denote the collection of information sets at which Player i moves. That is, $H_i = \{h : i(h) = i\}$. Then the set of actions for Player i is given by $A_i = \bigcup_{h \in H_i} A(h)$.

Now we can define a "pure strategy" for player i as a function s_i mapping from H_i into A_i with the property that $s_i(h) \in A(h)$ for all $h \in H_i$. Let S_i denote the set of pure strategies for Player i, and define $S = S_1 \times S_2 \times \ldots \times S_N$.

Nature's strategy is to move randomly according to a known distribution, as mentioned above.

We can think of a pure strategy as a complete contingent plan of actions. A pure strategy specifies a player's choice of action at each of her information sets, even for sets that are not reached. One way to interpret that requirement is that a strategy represents both what the player intends to do and what other players expect her to do.

<u>Example:</u>

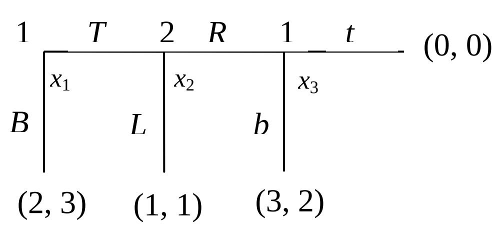

$S_1 = \{Tt, Tb, Bt, Bb\}$, $S_2 = \{L, R\}$.

Note that B is not a strategy for Player 1. A strategy must specify Player 1's choice at x_3, even though if he plays B at x_1 then node x_3 will never be reached.

A pure strategy profile s, together with the probabilities over Nature's moves, generates a probability distribution over terminal nodes. Let $o : S \to \Delta(Z)$ denote the "outcome" of a strategy profile, so that $o(s)[z]$ is the probability that terminal node z is reached when strategy profile s is played.

Now we can straightforwardly define expected payoffs as a function of strategy profiles, $U_i : S \to \mathbf{R}$: $U_i(s) = \sum_{z \in Z} o(s)[z]u_i(z)$.

A "mixed strategy" $\sigma_i \in \Delta(S_i)$ for Player i is a probability distribution over pure strategies. In the example on the previous page, $\{\sigma_1(Tt)= \sigma_1(Tb) = \sigma_1(Bb) = 1/3\}$ is a mixed strategy for Player 1.

A "behavior strategy $b_i \in \times_{h \in H_i} \Delta(A(h))$ is a mixed strategy such that the probability distributions over actions are independent across information sets. The mixed strategy $\{\sigma_1(Tt) = \sigma_1(Tb) = \sigma_1(Bb) = 1/3\}$ is not a behavior strategy, because $\sigma_1(t \mid T) = ½$, while $\sigma_1(t \mid B) = 0$.

Behavior strategies are easier to work with than unrestricted mixed strategies, so it would be nice if we could without loss of generality worry only about behavior strategies. Without going into detail, it turns out that we can!

Proposition 3.a.1. (Kuhn): For games with perfect recall, anything that we can do with mixed strategies, we can do with behavior strategies.

3.a.iv. Corresponding strategic forms and Nash equilibrium

Given an extensive form game G, consider the set of players 1, … N, the set of strategies for each player S_i, and the maps from strategy profiles to expected payoffs $U_i : S \rightarrow \mathbf{R}$. Those three objects define a normal form game, which is the "corresponding strategic form game" of the original extensive form game.

<u>Example:</u> Baby centipede game.

The extensive form is

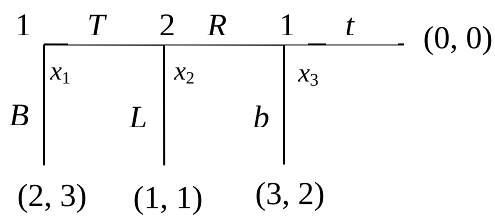

The corresponding normal form is

		Player 2	
		L	*R*
Player 1	*Tt*	1, 1	0, 0
	Tb	1, 1	3, 2
	Bt	2, 3	2, 3
	Bb	2, 3	2, 3

Note that the following extensive form game has the same corresponding normal form:

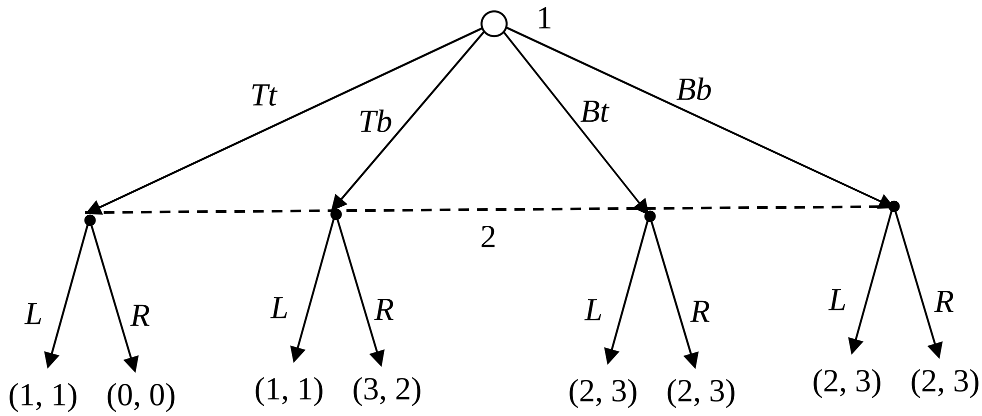

In general, any given normal form game corresponds to many different extensive form games. That is, there is more information in the extensive form than there is in the normal form.

Definition: A "Nash equilibrium" of an extensive form game is a Nash equilibrium of the corresponding strategic form game.

Nash equilibrium is a normal form concept, so the corresponding normal form of an extensive form game has all the information we need. To think about refinements of Nash equilibrium like backwards induction and subgame perfection, we need the extra information in the extensive form.

3.b. Refinements of Nash equilibrium

3.b.i. Backwards induction and subgame perfection

3.b.i.1. Backwards induction

Example: A Parent is driving to Disneyland with a Child in the back seat. The Child is making a lot of noise. The Parent says, "Pipe down back there, or I'll turn this car around." That situation is represented in the following extensive form game, where the Child's payoffs are given first:

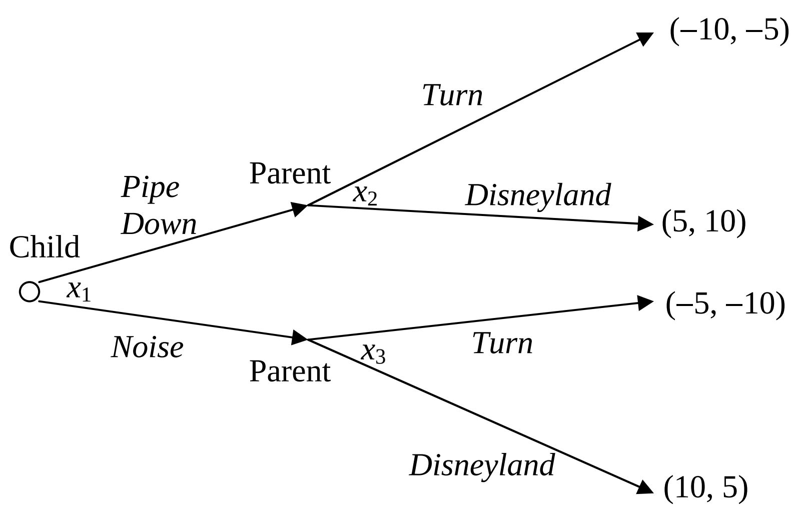

The strategy profile (*Pipe Down*, (x_2: *Disneyland*, x_3: *Turn*)) is a Nash equilibrium. But is the Parent actually willing to *Turn* at node x_3? Is that a "credible" threat?

Conditional on reaching node x_3, Parent's best response is *Disneyland*. The same is true at node x_2. Anticipating those decisions, Child knows that if he chooses *Pipe Down* he will end up with a payoff of 5, and if he chooses *Noise* he will get 10. Thus, he will choose *Noise*.

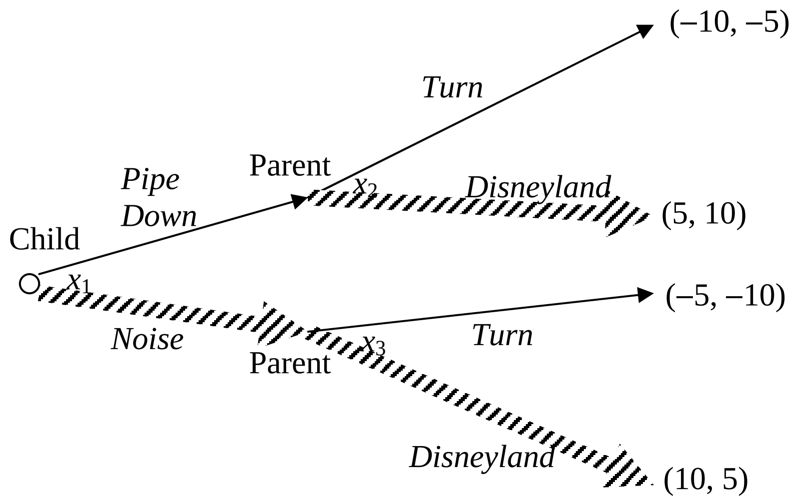

The strategy profile (*Noise*, (x_2: *Disneyland*, x_3: *Disneyland*)) is the "backwards induction solution" to the game. Note that that profile is also a Nash equilibrium, which is not a coincidence.

The following theorems are reasonably straightforward to prove:

Proposition 3.b.1: Every backwards induction solution of an extensive form game is a Nash equilibrium. (The converse is not true.)

Proposition 3.b.2: Every finite perfect information game has at least one backwards induction solution in pure strategies. Thus, every such game has a pure strategy Nash equilibrium.

There may be more than one backwards induction solutions if there are "ties" in payoffs. In that case, some solutions may involve mixed strategies.

Example: Baby centipede game.

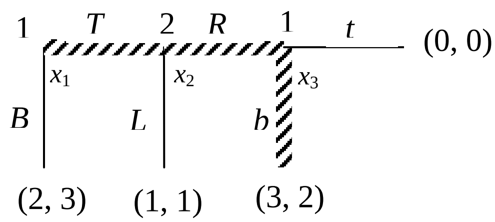

The backwards induction solution is (*Tb*, *R*).

Example: Backwards induction solutions may involve weakly dominated strategies. Consider the following game:

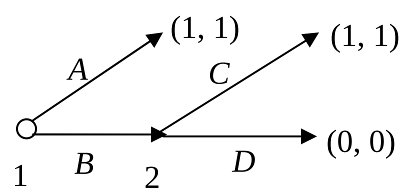

Even though *B* is weakly dominated for Player 1, (*B*, *C*) is a backwards induction solution.

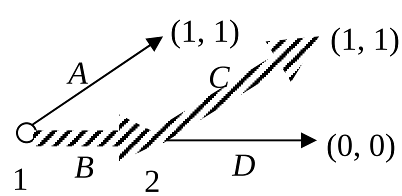

Note that (*A*, *C*) is also a backwards induction solution. In fact, the set of backwards induction solutions is $\{(\sigma_1(A) \in [0, 1], C)\}$.

3.b.i.2. Subgame perfection

Backwards induction does not work when there are non-trivial information sets:

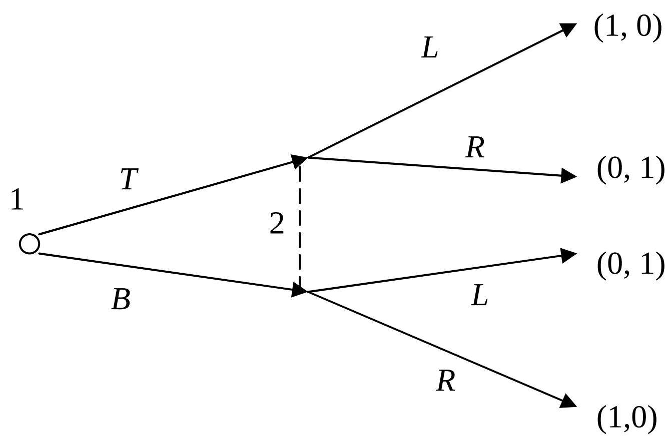

Subgame perfection is a related idea. As with backwards induction, the idea of subgame perfection is to require rationality even off the path of play.

Definition: A "subgame" is a subtree that i) starts at a decision node, and ii) contains no broken information sets. That is, if an information set contains a node in the subgame, then every node in the information set is in the subgame.

Example:

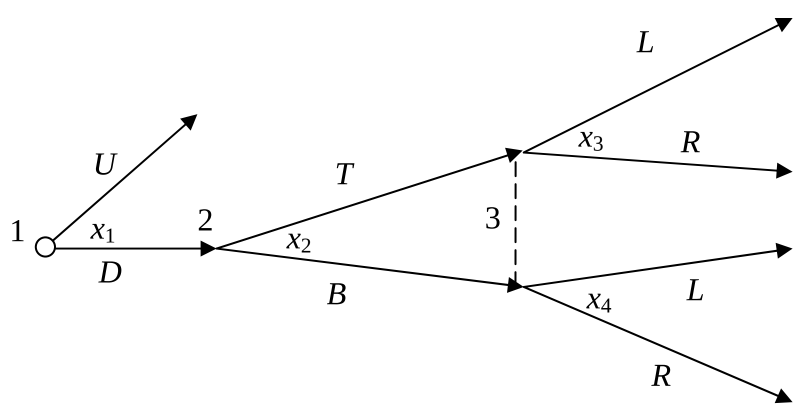

The subtrees starting at nodes x_1 and x_2 are subgames. (The one starting at x_2 is a "proper subgame.") The subtrees starting at x_3 and x_4 are not subgames.

Definition: A "subgame perfect equilibrium" is a profile of behavior strategies such that their restriction to any subgame forms a Nash equilibrium of that subgame.

In games of perfect information, the set of subgame perfect equilibria and the set of backwards induction solutions are the same. The advantage of subgame perfection is that it is defined even for games with imperfect information or infinite horizons.

Every subgame perfect equilibrium is a Nash equilibrium. If a game has no proper subgames, then every Nash equilibrium is subgame perfect.

To find subgame perfect equilibria, the procedure is similar to backwards induction: replace subgames at the end of the game with their equilibrium payoffs, and repeat until you reach the initial node.

<u>Example:</u>

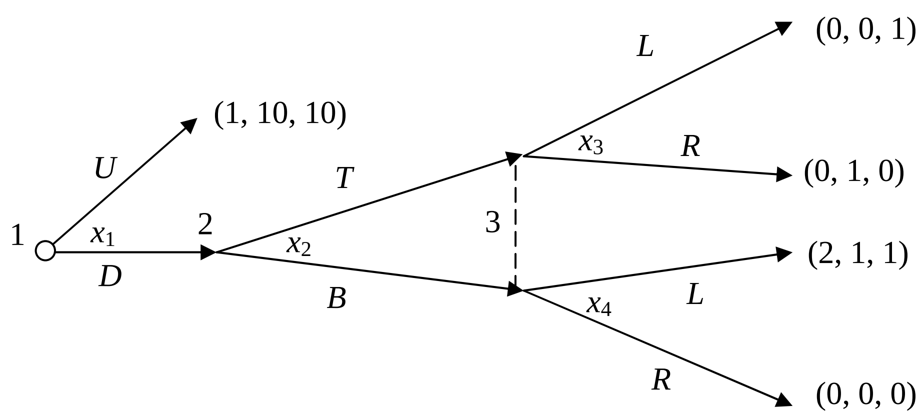

The unique subgame perfect equilibrium is (*D*, *B*, *L*). Strategy profiles (*U*, *T*, *L*) and (*U*, *B*, *R*), for example, are Nash equilibria but not subgame perfect.

Example: Centipede game.

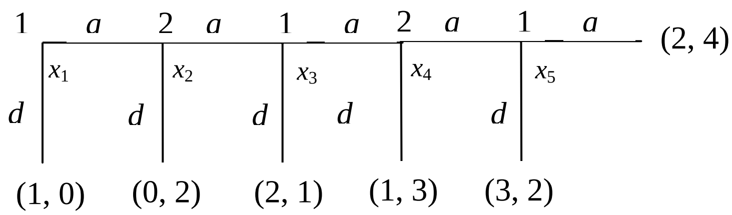

The unique subgame perfect equilibrium is (*ddd*, *dd*). More generally, we can make the centipede game as long as we want:

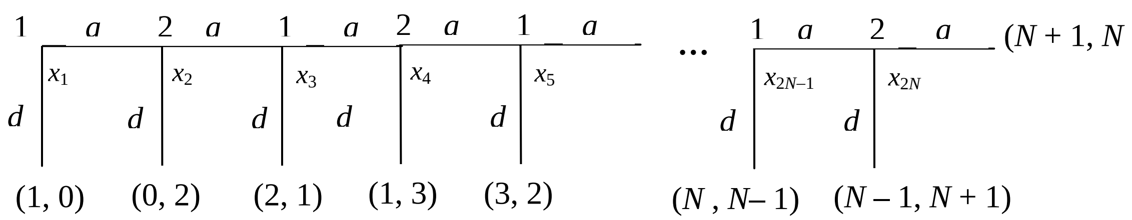

Still, the only subgame perfect equilibrium is (*ddd* … , *ddd* …). But does it make sense for Player 2 to choose *d* at node x_2? Given that Player 1 has behaved "irrationally" once by choosing *a* at node x_1, might he not do so again at x_3? Then Player 2 might get a payoff of 3 instead of 2. What if N = 1,000,000, and play reaches node $x_{500,000}$? Given that Player 1 has behaved "irrationally" 250,000 in a row, might he not do so again? By choosing *a* at that point, Player 2 gets a potential gain of 250,000 at a risk of only 1. How would you play the centipede game?

Subgame perfection requires a lot of rationality, especially in long games.

The idea of subgame perfection is to require common knowledge of rationality even off the equilibrium path of play. But if a player finds herself at a node off the equilibrium path, then she knows that someone has behaved irrationally, by deviating, so how can rationality still be common knowledge? The solution is for players to treat deviations as one-time events that will not be repeated.

3.b.ii. Problems with subgame perfection

Aside from the issues described above, it turns out that the set of subgame perfect equilibria can depend on what seem like irrelevant transformations of the game.

Example:

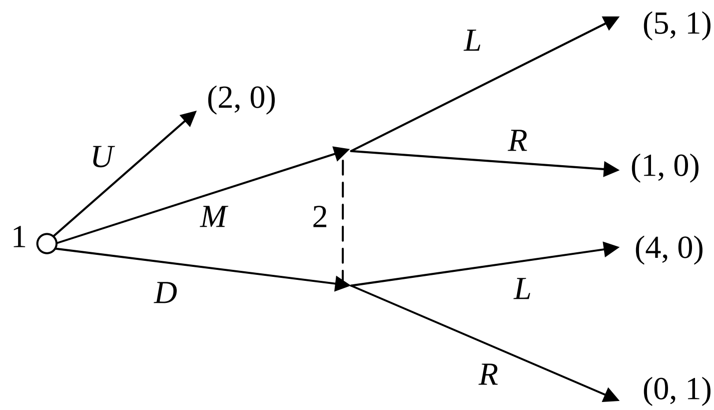

In the above game, (U, R) is a Nash equilibrium. It is also subgame perfect. (Why?)

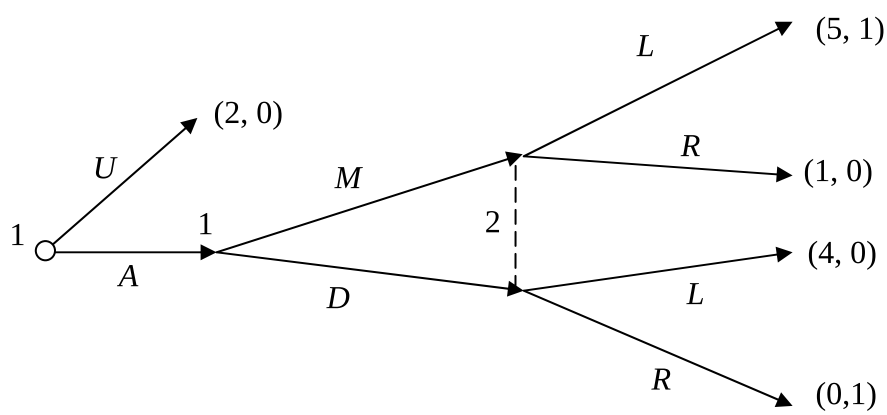

In this game, the equivalent (in some sense) strategy profiles (UM, R) and (UD, R) are Nash equilibria, but they are not subgame perfect. (AM, L) is the only subgame perfect equilibrium.

We will see other problems with subgame perfection in Section 3.d.

3.c. Repeated games

<u>Example:</u> Let's add a couple of actions to the Battle of the Sexes:

		Woman		
		Fight	*Ballet*	C_2
	Fight	2, 1	0, 0	6, 0
Man	*Ballet*	0, 0	1, 2	0, 0
	C_1	0, 0	0, 0	5, 5

Even with the additions, the only Nash equilibria are (*Fight*, *Fight*), (*Ballet*, *Ballet*), and ($\sigma_{Man}(Fight) = 2/3$, $\sigma_{Woman}(Fight) = 1/3$). It is sad that ($C_1$, C_2) is not an equilibrium, because then both players could get a payoff of 5. Unfortunately, Player 1 would want to deviate to *Fight*.

But consider the game $G(2)$, where the above game is played twice: first they play the game once, then both players' actions are revealed to each other, and then they play the game again. The payoffs from $G(2)$ are the sum of the payoffs in the two stages.

There exists a subgame perfect equilibrium in which (C_1, C_2) is played in period 1. A pure strategy s_i for Player i is a pair of functions, $s_i^1 \in \{Fight, Ballet, C_i\}$ and $s_i^2 : \{Fight, Ballet, C_1\} \times \{Fight, Ballet, C_2\} \to \{Fight, Ballet, C_i\}$. (Note that each player has $3 \times 3^9 = 3^{10} = 59{,}049$ pure strategies in $G(2)$.)

The following strategies form a subgame perfect equilibrium of $G(2)$:

$$s_i^1 = C_i,\ s_i^2(a_1, a_2) = \begin{cases} Fight & \text{if } a_1 = C_1 \\ Ballet & \text{otherwise.} \end{cases}$$

That is, if Man does not deviate from C_1 in period 1, then he is rewarded with his favorite stage-game equilibrium, (*Fight*, *Fight*) in period 2. Otherwise, (*Ballet*, *Ballet*) is played.

First, note that this strategy profile is a Nash equilibrium. If Man follows the equilibrium, he gets a payoff of 5 + 2 = 7. If he deviates in the first round, he gets at most 6 + 1 = 7, so deviating is not profitable. In the second round, he has no incentive to deviate. Woman has no incentive to deviate in either period – the equilibrium calls for her to play a static best response in both periods, and her action in period 1 does not affect play in period 2. So both players are playing best responses.

The equilibrium is subgame perfect because both (*Fight*, *Fight*) and (*Ballet*, *Ballet*) are Nash equilibria of the stage game: each of the nine proper subgames is that stage game.

That example suggests that our predictions from one-shot normal form games may change when the game is repeated. Many real-world situations, like price competition between firms or being interrogated by the police after robbing a bank with your buddy, seem to have the flavor of repeated interaction. Let us consider repeated games, which are a special kind of extensive form game, more generally.

3.c.i. Definitions and discounting

Remember that a normal form game G consists of three things: a set of N players, a set of actions A_i for each player, and a payoff function $u_i : A \to \mathbf{R}$ for each player. We will define the "repeated game" $G(T, \delta)$ as follows:

- The "stages" of the game are $t = 1, 2, \dots, T$, where T may be infinite.
- A "history" $h^t = (a^1, a^2, \dots, a^{t-1})$ is a list of the action profiles $a^s \in A$ played in each period $s < t$. The history h^t is commonly known at the start of period t. Let H^t denote the set of possible period-t histories, A^{t-1}. For the sake of notational consistency, we will think of $H^1 = \{h^1\}$ as a "dummy" singleton history for period 1. H^{T+1} is the set of complete histories of the game, which correspond to terminal nodes.
- A "pure strategy" for Player i is $s_i = (s_i^1, s_i^2, \dots, s_i^T)$, where $s_i^t : H^t \to A_i$. After observing history h^t, Player i chooses action $s_i^t(h^t)$. The set of pure strategies for Player i is S_i.
- A "behavior strategy" for Player i is $\sigma_i = (\sigma_i^1, \sigma_i^2, \dots, \sigma_i^T)$, where $\sigma_i^t : H^t \to \Delta(A_i)$. After observing history h^t, Player i chooses action a_i with probability $\sigma_i^t(h^t)[a_i]$. The set of behavior strategies for Player i is Σ_i.
- The "outcome" of a pure strategy profile s is a history $h^{T+1}(s)$ defined recursively as $a^1 = s^1(h^1)$, and $a^t = s^t(a^1, \dots, a^{t-1})$ for $2 \le t \le T$. Similarly, the outcome of a behavior strategy profile σ is a probability distribution $P(\sigma) \in \Delta(H^{T+1})$.
- The value $\delta \in (0, 1)$ is the "discount factor." For simplicity, we will assume that it is the same for all players.

Now all we need to complete our description of the game is payoffs. The payoff to Player i from a complete history h^{T+1} is the sum of Player i's payoffs in each stage, exponentially weighted by the discount factor:

$$\hat{U}_i(h^{T+1}) = u_i(a^1) + \delta u_i(a^2) + \dots + \delta^{T-1} u_i(a^T).$$

It will sometimes be convenient to rescale those payoffs so that they are directly comparable to the stage game payoffs. If we define

$$\hat{U}_i(h^{T+1}) = \frac{1-\delta}{1-\delta^T}[u_i(a^1) + \delta u_i(a^2) + \dots + \delta^{T-1} u_i(a^t)],$$

then the set of possible vectors of utilities are constant (and equal to the set for the stage game) as δ and T vary. When T is infinite, define payoffs as

$$\hat{U}_i(h^{\infty}) = (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}u_i(a^t).$$

With the rescaling, note for example that if $h^{T+1} = (a, a, \dots, a)$, then $U_i(h^{T+1}) = u_i(a)$.

Finally, we can write payoffs as functions of strategy profiles. Given a pure strategy profile s, player i's payoff is given by $U_i(s) = \hat{U}_i(h^{T+1}(s))$. The expected payoff from a mixed strategy is defined in the usual way.

3.c.ii. Finitely repeated games

Example: The Prisoners' Dilemma, repeated $T < \infty$ times.

		Player 2	
		C	D
Player 1	C	$-1, -1$	$-4, 0$
	D	$0, -4$	$-3, -3$

Remember that the only Nash equilibrium of the stage game is (D, D). In any subgame perfect equilibrium of $G(T, \delta)$, therefore, (D, D) must be played in period T. That means that what happens in the second-to-last period cannot affect play in the last period. So players must choose static best responses in period $T - 1$, which means an equilibrium of the stage game (that is, (D, D)) will be played. By similar reasoning, (D, D) will be played in every period.

Proposition 3.c.1: Suppose that stage game G has a unique Nash equilibrium α^*, and that $T < \infty$. Then the only subgame perfect equilibrium of $G(T, \delta)$ is to play α^* in every period.

If the stage has more than one Nash equilibrium, then the finitely repeated game has more than one subgame perfect equilibrium. For example, for each t we could specify a stage-game equilibrium to be played in that period, regardless of history. That profile is a subgame perfect equilibrium.

Proposition 3.c.2: Let $\alpha^*(t)$ map from periods 1, … , T into the set of stage-game Nash equilibria. Then the strategy profile defined by

$$\sigma_i^t(h^t) = \alpha_i^*(t) \text{ for each period } t, \text{ each history } h^t \in H^t, \text{ and each player } i$$

is a subgame perfect equilibrium of $G(T, \delta)$.

3.c.iii. Infinitely repeated games

As a preliminary, let's define "minmax" strategies and payoffs.

Given a normal form stage game, let α_{-i} be a mixed action profile for all players other than player i. Let $w_i(\alpha_{-i}) = \max_{a_i \in A_i} u_i(a_i, \alpha_{-i})$. That is, $w_i(\alpha_{-i})$ is Player i's payoff from best-responding to α_{-i}.

The "minmax payoff" $\underline{v}_i$ for Player i is defined as

$$\underline{v}_i = \min_{\alpha_{-i}} w_i(\alpha_{-i}) = \min_{\alpha_{-i}} \left[\max_{a_i \in A_i} u_i(a_i, \alpha_{-i}) \right].$$

That is, the other players are "minimizing" the value of Player i's "maximum" payoff, when he best responds. The "minmax strategy" that they use against Player i is

$$m_i = \arg\min_{\alpha_{-i}} w_i(\alpha_{-i}).$$

Proposition 3.c.3: The payoff to Player i in any Nash equilibrium of any repeated game (finite or infinite) is at least $\underline{v}_i$ (average per period).

Proof idea: Player i can guarantee herself no less than $\underline{v}_i$ just by playing a static best response in every period. Thus, she has a profitable deviation from any strategy profile that gives her a lower payoff than that.

Example:

		Player 2: C	D
Player 1	A	2, 1	0, 1
	B	3, 3	–1, 2

To find m_2, note that for Player 2 C weakly dominates D. So to minmax Player 2, Player 1 just chooses the action (A) that minimizes Player 2's payoff from playing C. To find m_1, note that Player 1's lowest possible payoff when Player 2 plays C is greater than his highest possible payoff when Player 2 plays D. So to minmax Player 1, Player 2 chooses D. That is, $m_1 = D$, $\underline{v}_1 = 0$, $m_2 = A$, and $\underline{v}_2 = 1$.

Example: Matching pennies.

		Player 2: *Heads*	*Tails*
Player 1	*Heads*	1, –1	–1, 1
	Tails	–1, 1	1, –1

To find m_1 and $\underline{v}_1$, observe that if Player 2 chooses *Heads* with probability p, then $w_1(p) = \max\{2p - 1, 1 - 2p\}$. Since $w_1(p) \geq 0$ for all $p \in [0, 1]$, with strict inequality for $p \neq 0.5$, $m_1 = \sigma_2(Heads) = 0.5$, and $\underline{v}_1 = 0$. Similarly, $m_2 = \sigma_1(Heads) = 0.5$, and $\underline{v}_2 = 0$.

3.c.iii.1. Folk theorems

Proposition 3.c.3 tells us that any Nash equilibrium must give each player a payoff of at least $\underline{v}_i$. A folk theorem, roughly, tells us that if players are patient enough (that is, if they care enough about the future), then that lower bound is the only restriction on equilibrium payoffs. That is, any feasible payoff vector that gives each player at least her minmax payoff is achievable in equilibrium.

Let V be the set of feasible payoffs in a stage game G. That is,

$$V = \text{co}\{\ (u_1(a), \ldots\ u_N(a)) : a \in A\}.$$

With the appropriate scaling, V is also the set of feasible payoffs in the infinitely repeated game $G(\infty, \delta)$.

Here is one of the most general folk theorems:

Proposition 3.c.4 (Fudenberg and Maskin's subgame perfect folk theorem): Suppose that the dimension of V is N. (That is, V has a nonempty interior.) Choose any vector of payoffs $v \in V$ such that $v_i > \underline{v}_i$ for every player i. Then there exists $\underline{\delta} \in (0, 1)$ such that whenever $\delta \in (\underline{\delta}, 1)$, there exists a subgame perfect equilibrium σ^* of the infinitely repeated game $G(\infty, \delta)$ such that $U(\sigma^*) = v$.

Proof intuition: When $v_i > \underline{v}_i$ for every player i, there is always a feasible, lower payoff to punish a deviating player with. The hard part is to make sure that punishing a deviator is an equilibrium – that is, to make the punishers want to punish, which can be tricky if the strategy profile that minmaxes Player 1 also gives Player 2 a low payoff.

3.c.iii.2. The principle of optimality

The principle of optimality (also known as the one-shot deviation principle) is most useful when we deal with infinitely repeated games, but it also holds for finitely repeated games and, more generally, for any multistage game in which players move simultaneously within a period and all past actions are publicly observed.

Proposition 3.c.5 (Principle of optimality, one-shot deviation principle): A strategy profile σ is a subgame perfect equilibrium if and only if no player at any history (on or off the equilibrium path) can increase her utility from that point on by deviating from σ exactly once (at that history) and playing σ again afterward.

(Technically, a condition on the utility functions called "continuity at infinity" is required for the principle of optimality to hold. Roughly, continuity at infinity means that what happens in the far future becomes vanishingly unimportant. That condition obviously holds for finite games, and it also holds for infinitely repeated games with discounting.)

The principle of optimality makes checking for subgame perfection much easier by reducing the number of deviations that we need to check. The idea is that if there is *any* profitable deviation, then there must be a profitable one-shot deviation.

Example: The expanded Battle of the Sexes:

		Woman *Fight*	Woman *Ballet*	Woman C_2
Man	*Fight*	2, 1	0, 0	6, 0
Man	*Ballet*	0, 0	1, 2	0, 0
Man	C_1	0, 0	0, 0	5, 5

For what values of the discount factor does the following strategy profile constitute a subgame perfect equilibrium of the infinitely repeated game $G(\infty, \delta)$?

$$s_i^1 = C_i, \text{ and for } t > 1 \; s_i^t(h^t) = \begin{cases} C_i & \text{if } a_1^s = C_1 \text{ and } a_2^s = C_2 \;\forall s < t \\ Ballet & \text{otherwise.} \end{cases}$$

That is, (C_1, C_2) is played in every period until someone deviates, and then (*Ballet*, *Ballet*) is played forever. There are two types of histories that we need to check: those when no one has deviated, and those after a deviation.

If no one has deviated, Man's equilibrium continuation payoff is $\frac{5}{1-\delta}$. Since any deviation results in the same future play, Man's best deviation is his best static deviation: *Fight*. That deviation yields a continuation payoff of $6+\delta\frac{1}{1-\delta}$. Woman's equilibrium continuation payoff is also $\frac{5}{1-\delta}$. Her best deviation is to either *Fight* or *Ballet*, both of which yield continuation payoff $0+\delta\frac{2}{1-\delta}$.

In any history after a deviation, Man's equilibrium continuation payoff is $\frac{1}{1-\delta}$. Now any further deviation does not affect future play, so again Man's best deviation is his best static deviation: either *Fight* or C_1. That deviation yields a continuation payoff of $0+\delta\frac{1}{1-\delta}$. Woman's equilibrium continuation payoff is $\frac{2}{1-\delta}$. Her best deviation is either *Fight* or C_2, both of which yield continuation payoff $0+\delta\frac{2}{1-\delta}$.

So the strategies are subgame perfect if the following four inequalities hold:

$$\frac{5}{1-\delta} \geq 6+\frac{\delta}{1-\delta}, \; \frac{5}{1-\delta} \geq \frac{2\delta}{1-\delta}, \; \frac{1}{1-\delta} \geq \frac{\delta}{1-\delta}, \text{ and } \frac{2}{1-\delta} \geq \frac{2\delta}{1-\delta}.$$

The last three hold for any $\delta \in (0, 1)$, so the condition for subgame perfection is just

$$\frac{5}{1-\delta} \geq 6+\frac{\delta}{1-\delta}, \text{ or } \delta \geq \frac{1}{5}.$$

3.d. Refinements of subgame perfection

3.d.i. Motivating examples

<u>Example 1:</u>

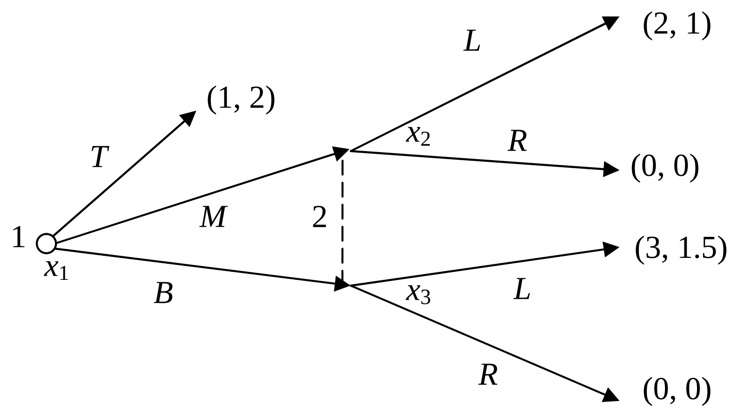

The strategy profile (T, R) is a Nash equilibrium, and it is subgame perfect.

However, no beliefs that Player 2 can have about which decision node she's at in her information set are consistent with R giving higher expected utility than L. The action R is *conditionally* strictly dominated.

We ought to require that at nontrivial information sets, players maximized expected utility according to some beliefs about which node they're at (and about future play by other players.) Those beliefs should be consistent with Bayes' rule applied to other players' strategies whenever possible. In this example, it is not possible. (Why?)

Example 2: Selten's horse

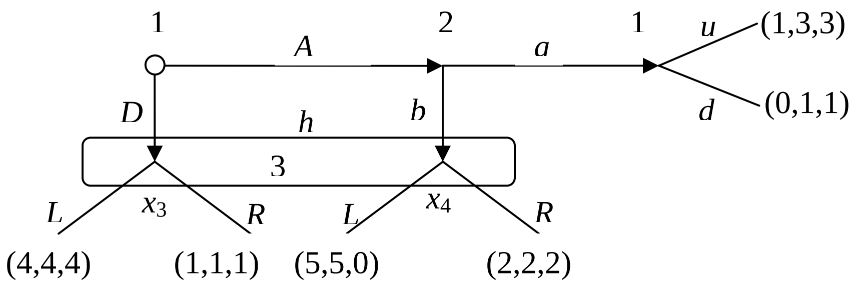

Strategy profile (*Du*, *a*, *L*) is a Nash equilibrium. It is also subgame perfect. However, Player 2's choosing *a* is inconsistent with knowledge of Player 3's equilibrium strategy.

Example 3:

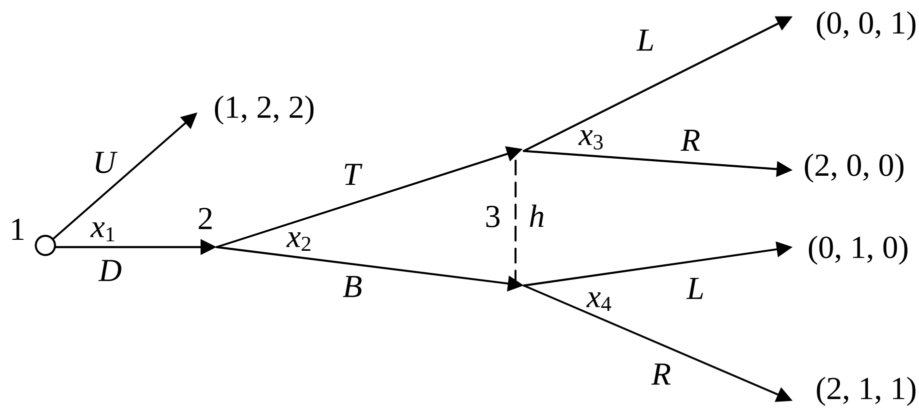

The unique subgame perfect equilibrium is (*D*, *B*, *R*), which is fine. However, the profile (U, $\sigma_2(B) = 2/3$, L) is also a Nash equilibrium. There is a belief over which nodes he's at that Player 3 could have that would lead him to pick *L* (that is, high probability on x_3), but that belief is in some sense inconsistent with knowledge of Player 2's strategy.

When Player 1 chooses *U*, Bayes' rule does not give a conditional probability $\mu(x_3 \mid h)$ since information set *h* is reached with probability zero. However, we *can* apply Bayes' rule to Player 2's strategy, assuming that play reaches x_2 (which it must if it reaches *h*). In that case, $\mu(x_3 \mid h) = 1/3$, and Player 3's best response is *R*.

3.d.ii. Perfect Bayesian equilibrium

First, let's be more formal about beliefs at information sets. Let decision node x be an element of Player i's information set h. Then "conditional belief" $\mu_i(x \mid h)$ is the probability of being at node x given that Player i is at some node in information set h: $\sum_{x \in h} \mu(x \mid h) = 1$.

Given Nature's moves and a strategy profile σ, Bayes' rule pins down beliefs at all information sets that are reached with positive probability when σ is played.

A few more definitions:

- An "assessment" is a (behavior) strategy / conditional belief pair (σ, μ), where the function $\mu: X \rightarrow [0, 1]$ gives conditional beliefs at each information set.

- An assessment (σ, μ) is "sequentially rational" if playing σ_i maximizes expected utility given μ for each player i at each of Player i's information sets. That is, σ is a best response for all players given μ.

- An assessment (σ, μ) is a "perfect Bayesian equilibrium" if
  - it is sequentially rational, and
  - beliefs μ are given by Bayes' rule applied to Nature's move and to σ "whenever possible."

Proposition 3.d.1: If (σ, μ) is a perfect Bayesian equilibrium, then σ is a subgame perfect equilibrium.

"Whenever possible" includes all information sets that are reached with positive probability when σ is played. It also includes some that are not, as in Example 3 in Section 3.d.i. In example 1, on the other hand, deriving beliefs from Bayes' rule is neither possible nor "possible." We could be more formal about defining "whenever possible," but we won't.

Note: MWG's "weak perfect Bayesian equilibrium" puts no restrictions on beliefs at information sets off the equilibrium path, and so is easier to define formally. However, a weak perfect Bayesian equilibrium need not be subgame perfect.

<u>Example 3:</u>

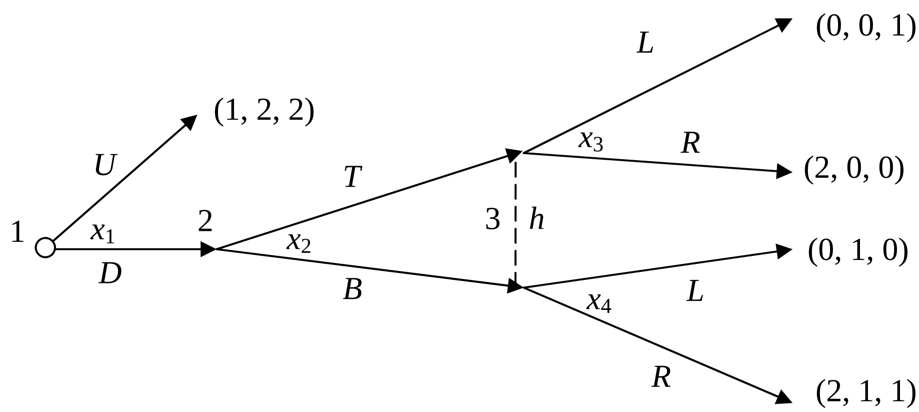

The unique subgame perfect equilibrium is (*D*, *B*, *R*), and so (*D*, *B*, *R*) must be the strategy part of any perfect Bayesian equilibrium. When (*D*, *B*, *R*) is played, every information set is reached, and so beliefs are given by Bayes' rule: $\mu_3(x_3 \mid h) = 0$. (Technically, we also need to describe beliefs at the trivial information sets, but those beliefs are also trivial.)

Example 2: Selten's horse

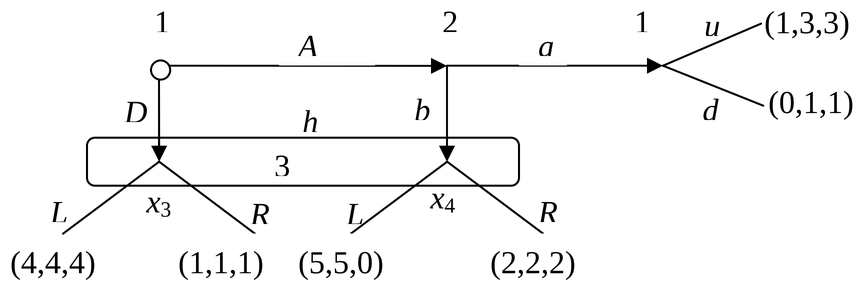

Homework exercise.

Example 1:

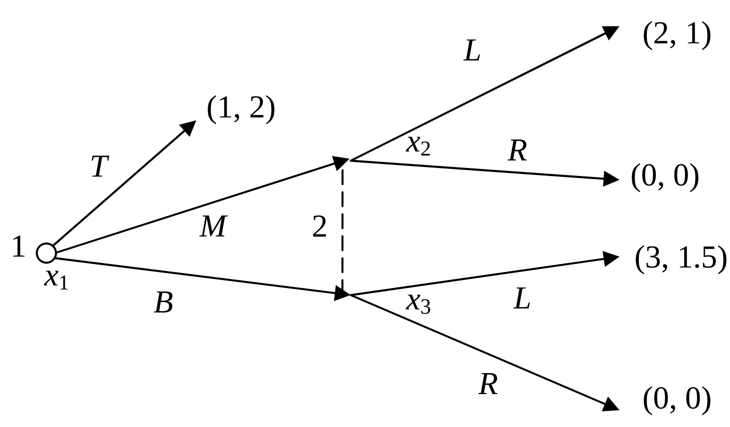

As we noted earlier, Player 2's conditional best response at her information set is *L*, regardless of her beliefs. So in any perfect Bayesian equilibrium, she plays *L*. Player 1's best response to *L* is *B*. So the unique perfect Bayesian equilibrium is the assessment $((B, L), \mu_2(x_2 \mid h) = 0)$.

3.d.iii. Sequential equilibrium

Even perfect Bayesian equilibrium has its problems:

<u>Example:</u>

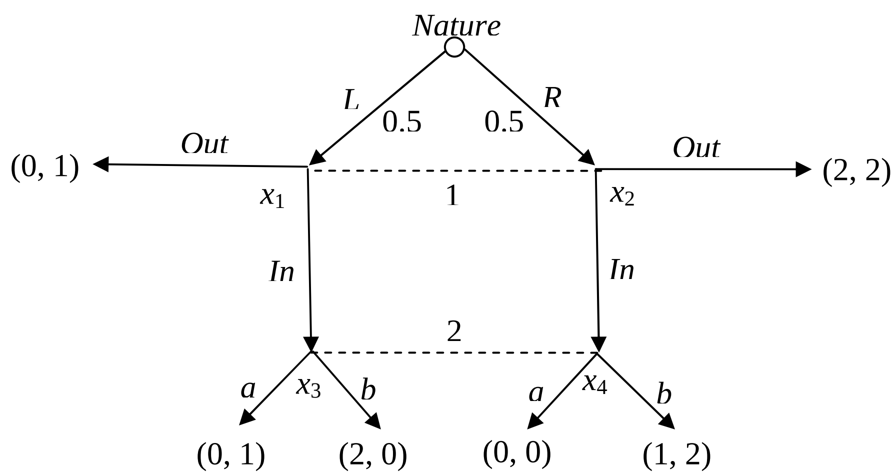

Strategy profiles (*In*, *b*) and (*Out*, *a*) are Nash equilibria. They are also subgame perfect. (Why?)

There are infinitely many pure strategy perfect Bayesian equilibria, but they fall into two categories:

1) $s_1 = In$, $\mu_1(x_1) = 0.5$, $s_2 = b$, $\mu_2(x_3) = 0.5$, and

2) $s_1 = Out$, $\mu_1(x_1) = 0.5$, $s_2 = a$, $\mu_2(x_3) \geq \frac{2}{3}$.

Equilibria in the second category are "bad." Since Player 1's actions cannot depend on Nature's move (which he does not observe), and Player 2 knows that, we ought to require that $\mu_2(x_3) = 0.5$, even though it is neither possible nor "possible" for Bayes' rule to determine Player 2's beliefs.

Definition: A strategy profile is "completely mixed" if it assigns strictly positive probability to all pure strategies (so that every information set is reached with positive probability).

Definition: An assessment (σ, μ) is a "sequential equilibrium" if

- it is sequentially rational, and
- there exists a sequence of completely mixed strategies $\{\sigma^k\}_{k=1}^{\infty}$ that induce beliefs $\{\mu^k\}_{k=1}^{\infty}$ (given by Bayes' rule) such that $\lim_{k\to\infty}\sigma^k = \sigma$ and $\lim_{k\to\infty}\mu^k = \mu$.

<u>Example:</u>

In the example above, the first perfect Bayesian equilibrium is also a sequential equilibrium, but those in the second category are not: for any completely mixed strategy profile σ^k, the induced belief $\mu^k(x_3) = 0.5$, so $\lim_{k\to\infty}\mu^k(x_3)$, if it exists, must equal 0.5.

To see that the first equilibrium is sequential, choose $\sigma_1^k(In) = \sigma_2^k(b) = 1 - \varepsilon^k$, where $\varepsilon \in (0, 1)$. Then $\mu^k(x_1) = \mu^k(x_3) = 0.5$ for all k, and $\lim_{k\to\infty}\sigma_1^k(In) = \lim_{k\to\infty}\sigma_2^k(b) = 1$.

Proposition 3.d.2: If (σ, μ) is a sequential equilibrium, then (σ, μ) is a perfect Bayesian equilibrium (and so σ is a subgame perfect equilibrium and a Nash equilibrium).