Sequential games

Moty Katzman

November 22, 2016

An example

Alice and Bob play the following game: Alice goes first and chooses A, B or C. If she chose A, the game ends and both get 0. If she chose B, Bob can either choose D resulting in utilities of 2 and 3, or he can choose E, resulting in utilities -2 and 2. If Alice chose C, Bob can either choose F resulting of utilities of 3 and 5, or G, resulting in utilities of -6 and 6. We can summarize this game with a rooted, directed tree as follows:

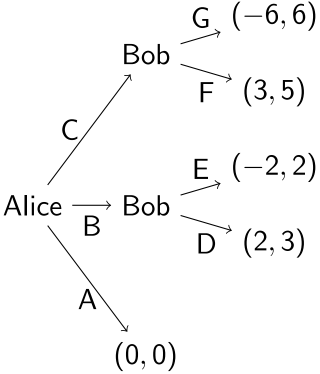

Example: Solving games using backward induction

If Bob has to choose between D and E, he will choose the outcome with highest utility for him, namely D, and if he needs to choose between F and G, he will choose G. So when Alice has to choose between A, B and C, she is choosing between utilities 0, 2 and -6, so she will choose B. The optimal strategies are for Alice to choose B and then for Bob to choose D, resulting in a payoff of 2 for Alice and 3 for Bob.
There is an outcome which is better for both Alice and Bob: If Alice chooses C and Bob F, Alice's payoff is $3 > 2$ and Bob's payoff is $5 > 3$. The problem is that if Alice chooses C, Bob will chose G for a payoff of $6 > 5$. But suppose we change the payoff $(-6, 6)$ to $(-6, 4)$: are we hurting Bob? No! Now the optimal strategy for both players leads to the optimal payoff $(3, 5)$.
Surprisingly, lowering one's payoffs in a sequential game can lead to a better outcome.

Directed rooted trees

Definition. A *rooted tree* is a directed graph with a distinguished vertex, the *root*, which is connected to every other vertex with a unique directed path.
Vertices with no arrows coming out from them are *leaves*.
Leaves will encode payoffs, and the other vertices will be labeled by the players of the game.
Edges coming out of a vertex will be labeled by the possible decisions at that stage of the game.
The *rank* of a rooted tree is the maximum of all lengths of directed paths in the tree.

Example: “The hungry lions” game

A pride of lions consist of a highly hierarchical group of lions. When prey is caught, the top lion eats first, followed by lion number two, and so on. Once an especially large and tasty prey was caught, so large and tasty that it would be impossible for lions to stop eating it, and they would fall asleep. If a lion falls asleep, he will be eaten by the lion immediately below him in the hierarchy, but not by any other lion. This lion will then also fall asleep. Should the top lion eat the prey?
Suppose there are n lions. If presented with the choice of eating the prey, should lion number n eat it? Yes, no one will bother to eat it. Should lion $n-1$ eat? No! It would fall asleep and be eaten. Should the $n-2$ lion eat? We know that lion $n-1$ won't eat it, so it is safe for it to eat and fall asleep, etc. So the top lion should eat if and only if n is odd!

Cournot duopoly revisited: the Stackelberg model

Two firms produce an identical product. Unlike in the Cournot model, the second firm will choose to produce q_2 *after* the first firm produces q_1. As before, the cost of production per unit is denoted c, and the price per unit of product is given by $p = a - b(q_1 + q_2)$. The second company's best response to company 1 producing q_1 units is $\mathrm{BR}_2(q_1) = (a-c)/2b - q_1/2$ resulting in total production of $q = q_1 + q_2 = q_1 + (a-c)/2b - q_1/2 = (a-c)/2b + q_1/2$. The first firm would then make a profit of

$$P(q_1) = (p-c)q_1 = \left(a - b\left(\frac{a-c}{2b} + q_1/2\right) - c\right) q_1 =$$

$$\left(\frac{a-c}{2} - \frac{q_1 b}{2}\right) q_1 = \frac{b}{2}\left(\frac{a-c}{b} - q_1\right) q_1$$

and $P(q_1)$ is maximized when $q_1 = (a-c)/2b$.

The second firm will produce $q_2 = \mathrm{BR}_2(q_1) = (a-c)/2b - q_1/2 = (a-c)/2b - (a-c)/4b = (a-c)/4b$, for a total production of $3(a-c)/4b$ which is higher than the Cournot Duopoly production of $2(a-c)/3b$.

The profit of the first firm is $\left(a - \frac{3}{4}(a-c) - c\right)\frac{a-c}{2b} = \frac{(a-c)^2}{8b}$ which is higher than Cournot profit of $(a-c)^2/9b$.

Moving first (=setting production levels first) gives an advantage.

Solving strictly competitive games with no chance moves

We now study more familiar games, similar to and chess and checkers.
In these games two players take turns making moves, until the game comes to an end and the outcome of the game is announced: either player I wins or player II wins or there is a draw.
(We shall call the player who moves first "player I" and the other player "player II".)
Definition. A *strictly competitive* game is a 2-player game in which for any two outcomes (u_1, u_2) and (u'_1, u'_2), $u_1 > u'_1$ implies $u_2 < u'_2$.
Since the preferences of player II in a strictly competitive game can be inferred from the preferences of player I, we can, and shall, keep track only of the outcomes of player I, e.g., the outcomes in a game of chess can be thought as being "player I wins", "the game is drawn", and "player I loses".

Subgames and subtrees

Definition. Let T be the tree associated with a sequential game G. For any vertex v in T we can construct a rooted subtree T' of T whose root is v, whose set of vertices V' consist of all vertices in T which can be reached by a directed path starting at v, and whose edges are all edges in T connecting vertices in V'. Such a subtree defines in turn a new game, the *subgame of G starting at v*.

Example: 2-pile Nim

There are two piles of stones, Alice and Bob take turns to make moves, each move consists of choosing a pile and taking way any number of stones from it. The person who takes the last stone wins.
Suppose that the piles have initially 5 and 8 stones and Alice goes first. Show that she has a strategy which ensures her victory.
Suppose that the piles have initially 7 stones each and Alice goes first. Show that Bob has a strategy which guarantees him victory. (Hint: analyze the game which starts with one stone in each pile.)

Zermelo's theorem

Definition. A sequential game is *finite* if its associated tree is finite.

Theorem.(Zermelo's theorem) Consider a finite, sequential game (with no chance moves) and let $\mathcal{S}$ the set of outcomes of the game. For any $S \subseteq \mathcal{S}$ either

(a) player I can force an outcome in S, or

(b) player II can force an outcome not in S.

Proof. Let T be the tree of this game; we proceed by induction on the rank r of T. If $r = 1$, T consist of a root and leaves only and either player I can choose a leaf with outcome in S or she is forced to choose a leaf with outcome not in S.

Assume now that $r > 1$ and that the result holds for all games with trees of rank less than r. Let $G_1, \ldots, G_n$ be the subgames resulting after player I makes her move, and let $T_1, \ldots, T_n$ be the corresponding trees. Note that the ranks of $T_1, \ldots, T_n$ are less than r, and we may apply the induction hypothesis to each of $G_1, \ldots, G_n$. We have two cases: either

(i) Player II can force an outcome not in S in each one of the games $G_1, \ldots, G_n$, or

(ii) for some G_i, player II cannot force an outcome not in S.

If (i) holds, conclusion (b) follows, while if (ii) holds player I has a strategy which starts with a move to game G_i which forces an outcome in S. $\square$

Strictly competitive, finite, sequential games have a value

Let the outcomes of such game be $u_1 < u_2 < \cdots < u_k$ where $<$ denotes player I's preference.
For any $1 \leq i \leq k$, let $W_i = \{u_i, \ldots, u_k\}$ and $L_i = \{u_1, \ldots, u_i\}$.
There exists a largest $1 \leq i \leq k$ for which player I can force an outcome in W_i;
player I cannot force an outcome in W_{i+1}, so player II can force an outcome in L_i.
We conclude that player I has a strategy that guarantees her at least u_i and player II has a strategy that guarantees him that player I will not do better than u_i.

Chess-like games

Corollary. In chess, either white can force a win, or black can force a win, or both sides can force at least a draw.

Proof. Here $\mathcal{S} = \{\text{white wins}, \text{draw}, \text{black wins}\}$. Apply Zermelo's Theorem to $S_1 = \{\text{white wins}\}$ and $S_2 = \{\text{white wins, draw}\}$ and deduce that

(a) either white can force a win or black can force a win or a draw, and

(b) either black can force a win or white can force a win or a draw.

So we either have a strategy which guarantees victory by one side, or, if this fails, both sides have a strategy that guarantees at least a draw. $\square$

Chomp!

A game of Chomp! starts with a rectangular array of pieces and two players alternate taking turns. A turn consists of a player choosing a piece and removing that piece and all other pieces above or to the right of the chosen piece. The player who takes the last piece loses the game.

Theorem. The player who goes first has a strategy which guarantees her victory.

We use the following *strategy-stealing argument*.

Proof. Assume that the Theorem fails; Zermelo's theorem implies that the second player has a strategy which guarantees him victory. If the first player chooses the upper-right piece, there is a choice of piece P for the second player which is the first step in a winning strategy. Let the first player then start with a choice of P, and let him follow the winning strategy with the roles reversed! □

Although it is known that there exists a winning strategy for Player I in Chomp!, it is not known what that strategy is.

(You can play this game in

`http://www.math.ucla.edu/ tom/Games/chomp.html`

Sequential games in strategic form

So far we described sequential games as rooted trees T where

(a) vertices which were not leaves were labeled with players,

(b) arrows originating from a vertex were labeled with actions available to the player corresponding to the vertex,

(c) arrows either pointed to leaves labeled with payoffs, or to other vertices as in (a)

Definition. A *strategy* for player i in the sequential game described by the rooted tree T is a function which takes any vertex v labeled i to an action labeled on an arrow originating from v.

Example

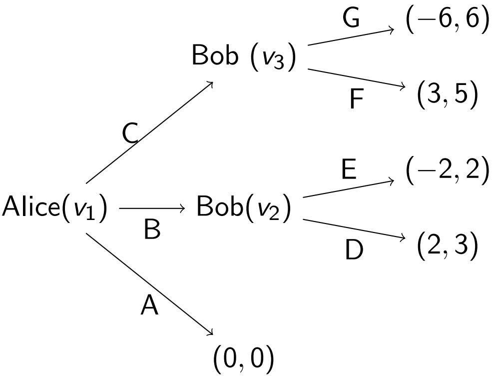

Alice has three strategies: $[v_1 \to A]$, $[v_1 \to B]$, and $[v_1 \to C]$. Bob has four strategies strategies $[v_2 \to D, v_3 \to F]$, $[v_2 \to D, v_3 \to G]$, $[v_2 \to E, v_3 \to F]$, $[v_2 \to E, v_3 \to G]$.

We can now give the strategic form of this game as follows

	$[D, F]$	$[D, G]$	$[E, F]$	$[E, G]$
A	0, 0	0, 0	0, 0	0, 0
B	2, 3	2, 3	-2, 2	-2, 2
C	3, 5	-6, 6	3, 5	-6, 6

Note that we have a pure-strategy Nash equilibrium $(B, [D, G])$, corresponding to our previous solution. There is also a new Nash equilibrium $(A, [E, G])$.

Monopolist vs. new entrant

Bob has a monopoly in the market of wireless widgets, and Alice is considering competing with Bob. If Alice does not enter the market, her payoff is zero, and Bob makes a three billion pound profit. If Alice enters the market, Bob can either fight her off by reducing prices so that he makes zero profit and Alice loses a billion pounds, or Bob can choose not to fight her and they both end up making a profit of a billion pounds. Here is the tree describing this sequential game

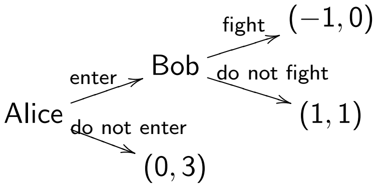

We use backward induction to solve this game an we see that the optimal strategies are for Alice to enter the market, and for Bob not to fight her off.

Consider the strategic form of this game

	fight	do not fight
enter	-1, 0	1, 1
do not enter	0, 3	0, 3

We have two Nash equilibria at (enter, do not fight) and at (do not enter, fight). The first corresponds to the solution above, what is the significance of the second? For Bob choosing to fight amounts to announcing that, no matter what Alice does, he will fight her off, but this is not credible as we know that it would be foolish for Bob to fight Alice if she enters the market. His strategic choice is *not credible*. The situation could be different if the game were to repeat, e.g., if Bob would have to face multiple potential competitors: now he may choose to fight to set an example to the others.

So in sequential games, strategic Nash equilibria do not necessarily correspond to feasible choice of strategy.

Subgame perfect Nash equilibria

Definition. Consider a sequential game with perfect information. A strategy profile which is a Nash equilibrium of the game is *subgame perfect* if its restriction to any subgame is also a Nash equilibrium.

Example: Market monopolist and new entrant, revisited.

Recall the game

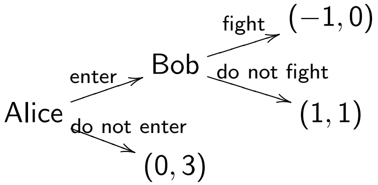

whose strategic form is

	fight	do not fight
enter	-1, 0	1, 1
do not enter	0, 3	0, 3

The game has two Nash equilibria: (*enter*, *do not fight*), (*do not enter*, *fight*). There is only one subgame to consider here, the one involving Bob's decision to fight or not, and clearly, *not fight* is the only Nash equilibrium. So (*enter*, *do not fight*) is a subgame perfect Nash equilibrium, and (*do not enter*, *fight*) is not.

Question from Mock exam 2013-14

The Klingons (a belicose alien civilization) invade the planet Romulus, and the Romulans need to decide whether to abandon their planet or to stay put. If they stay and the Klingons fight, both get a payoff of 0, whereas if the Klingons run away, Romulans get 2 and Kilngons get 1. If the Romulans run away, they get 1 and the Klingons get 2.

Immediately after landing on Romulus, the Klingon commander needs to decide whether to destroy her spaceships (and thus eliminating the option to run away if Romulans decide to stay put).

(a) Assuming everyone is rational and well informed, what should the Klingon commander do? Explain you reasoning in detail. (6 marks)

(b) Describe this situation in detail as a two-player game in strategic form. (6 marks)

(c) Find a pure strategy Nash equilibrium of this game which is not subgame perfect. (4 marks)

Solution

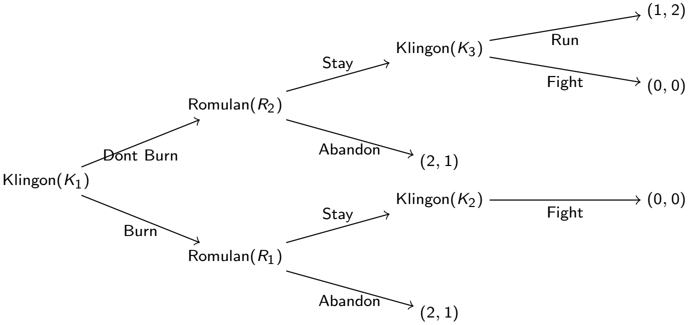

If Klingons do burn their spaceships, they will have to fight at K_2, and so Romulans would prefer to abandon at R_1 and Klingons will receive 2 units of utility. If they don't burn their spaceships, when facing a decision at K_3 they will decide to run, and so Romulans will decide to stay at R_2, resulting in 0 units of utility for the Klingons. So the Klingon general will order the spaceships burned.

Solution

Klingons' strategies amount to choosing action a at at K_1 and b at K_3 which we denote $[a, b]$. Romulans' strategies amount to choosing action c at R_1 and d at R_2 which we denote $[c, d]$. The game in strategic form is given by the following table, where Klingons are the row players and Romulans are the column players

	[*Abandon*, *Abandon*]	[*Abandon*, *Stay*]	[*Stay*, *Abandon*]	[*Stay*, *Stay*]
[*Burn*, *Run*]	2,1	2, 1	0, 0	0, 0
[*Burn*, *Fight*]	2,1	2, 1	0, 0	0, 0
[*DontBurn*, *Run*]	2,1	1, 2	2, 1	1, 2
[*DontBurn*, *Fight*]	2,1	0, 0	2, 1	0, 0

It is easy to verify that ([Dont Burn, Fight], [Abandon, Abandon]) is a Nash equilibrium, but it is not subgame perfect. To see this consider the subgame starting at K_3: this strategy call for the Klingons to fight, which is certainly not a Nash equilibrium as they gain by running instead.

Theorem. The backward induction solutions of a finite game of perfect information are subgame perfect Nash equilibria.
Proof. Let T be the tree of this game; we proceed by induction on the rank r of T. Call the player who moves at the root of T player I. If $r = 1$, T consist of a root and leaves only and player I will choose the strategy with highest payoff. This is clearly a subgame perfect Nash equilibrium.

Assume now that $r > 1$ and that the result holds for all games with trees of rank less than r. Let $G_1, \ldots, G_n$ be the subgames resulting after player I makes her move, and let $T_1, \ldots, T_n$ be the corresponding trees. Note that the ranks of $T_1, \ldots, T_n$ are less than r, and if we apply the induction hypothesis to each of $G_1, \ldots, G_n$ we obtain backward induction solutions of these which are subgame perfect Nash equilibria resulting in payoffs $u_1, \ldots, u_n$ for player I. Now player I needs to choose to move to one of the subgames $G_1, \ldots, G_n$, and she will choose the one with highest payoff for her. This results in a Nash equilibrium strategy. Note that a proper subgame of G is a subgame of one of $G_1, \ldots, G_n$ so the strategy is subgame perfect, because its restriction to $G_1, \ldots, G_n$ is subgame perfect and so its restriction to any subgame is a Nash equilibrium. $\square$

The material after this slide will not be covered in this year's MAS348

Imperfect information and information sets

We now deal with games in which some of the actions of players are not known to the other players: these are *games of imperfect information.*

Example. Any game in strategic form is a game of imperfect information: players do not know the other players actions until the game is finished.

Definition. Let G be a sequential game represented by a tree T. An *information set* for player i is a set of vertices V such that

(a) each vertex in V is labeled i,

(b) the set of arrows starting at each vertex in V have identical labels.

We partition all vertices which are not leaves into information sets, i.e., every such vertex is in precisely one information set.

A sequential game with partial information has as its strategies for player i the set of all functions which take an *information set* to an arrow starting at a vertex in that information set.

Imperfect information and information sets

Note that in a game of perfect information all information sets contain one vertex.
An information set of a sequential game is a set of vertices belonging to one player which are indistinguishable to that player.
Example. Matching pennies as a sequential game.

Example: Alice and Bob compete, Part I

Alice Ltd. is considering entering a new market in Freedonia dominated by Bob plc. If Alice enters the market, Bob can decide to resist or accommodate Alice's entrance. When Alice enters the market, she can choose to do so aggressively (e.g., lots of advertising) or not. If at the time of Alice's second decision, she knows Bob's choice of strategy, we have a sequential game described by the following tree.

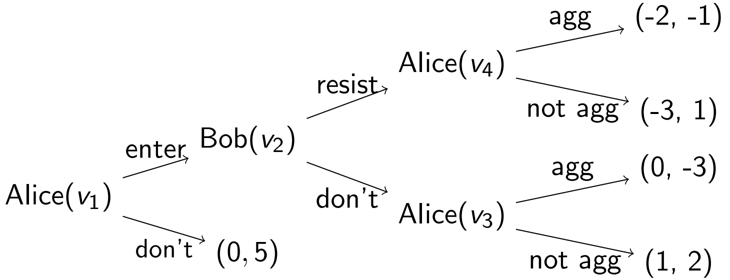

The backward induction solution of this game is straightforward: Alice will enter, Bob will not resist and Alice will *not* advertise

The normal form of this game is as follows:

	resist	don't resist
enter, aggressive at v_3, aggressive at v_4	(-2, -1)	(0, -3)
enter, not aggressive at v_3, aggressive at v_4	(-2, -1)	(1,2)
enter, aggressive at v_3, not aggressive at v_4	(-3, -1)	(0, -3)
enter, not aggressive at v_3, not aggressive at v_4	(-3, -1)	(1, 2)
don't enter, aggressive at v_3, aggressive at v_4	(0,5)	(0,5)
don't enter, not aggressive at v_3, aggressive at v_4	(0,5)	(0,5)
don't enter, aggressive at v_3, not aggressive at v_4	(0,5)	(0,5)
don't enter, not aggressive at v_3, not aggressive at v_4	(0,5)	(0,5)

Example: Alice and Bob compete, Part II

Assume now that when Alice enters the market she does not know whether Bob decided to resist or not; Alice makes her second decision without knowing whether she is in vertex v_3 or v_4. This implies that any of its strategies has the same values on v_3 or v_4, and we can express this formally by saying that $\{v_3, v_4\}$ is an information set of Alice.

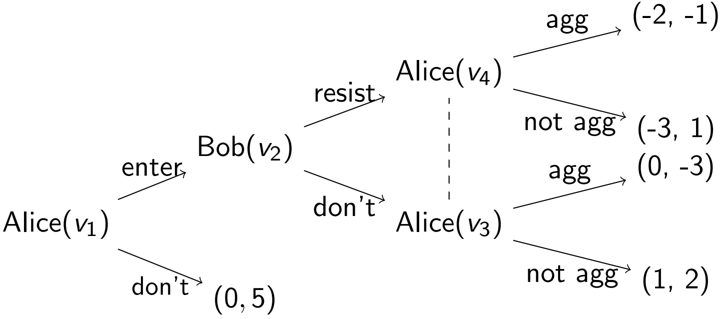

The normal from of this modified game is as follows:

	resist	don't resist
enter, aggressive at v_3, aggressive at v_4	(-2, -1)	(0, -3)
enter, not aggressive at v_3, not aggressive at v_4	(-3, -1)	(1, 2)
don't enter, aggressive at v_3, aggressive at v_4	(0,5)	(0,5)
don't enter, not aggressive at v_3, not aggressive at v_4	(0,5)	(0,5)

We cant solve this game with backward induction anymore, but we can find its Nash equilibria: any strategy profile where Alice stays out and Bob resists is a Nash equilibrium, and there is an additional one: ((enter, not aggressive at v_3, not aggressive at v_4) , dont resist).

An example

Consider the following game with perfect information.

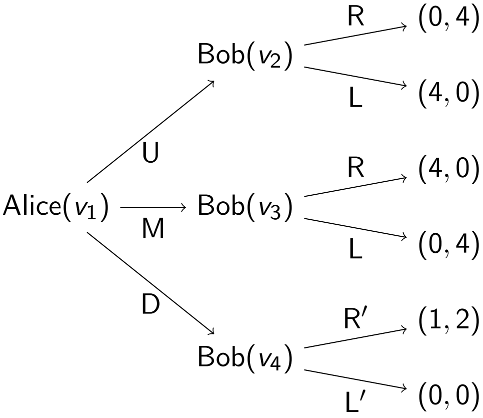

We easily solve this game using backward induction and obtain a solution $([D], [R'])$.

We now change the game and put vertices v_2 and v_3 in one information set: now Bob doesn't distinguish between these two vertices. Now Alice can mix strategies U and M with equal probabilities to produce an expected payoff of $2 > 1$. Formally, the game in strategic form is

	L L′	L R′	R L′	R R′
U	4, 0	4, 0	0, 4	0, 4
M	0, 4	0, 4	4, 0	4, 0
D	0, 0	1, 2	0, 0	1, 2

Alice's mixed strategy $(1/2, 1/2, 0)$ dominates strategy D and we reduce to the game

	L L′	L R′	R L′	R R′
U	4, 0	4, 0	0, 4	0, 4
M	0, 4	0, 4	4, 0	4, 0

	L L′	L R′	R L′	R R′
U	4, 0	4, 0	0, 4	0, 4
M	0, 4	0, 4	4, 0	4, 0

Now the first two and last two columns are indistinguishable, and if we ignore duplication we end up with the game

	L	R
U	4, 0	0, 4
M	0, 4	4, 0

which has mixed-strategy Nash equilibrium $((1/2, 1/2), (1/2, 1/2))$.

To define subgame perfect solutions of games without perfect information we need to modify the concept of a subgame.

Definition. Let G be a sequential game represented by a tree T. A *subgame* of G is a game corresponding to a subtree of T' of T with the property that every information set is entirely in T' or entirely outside T'.

Example

Consider the following game.

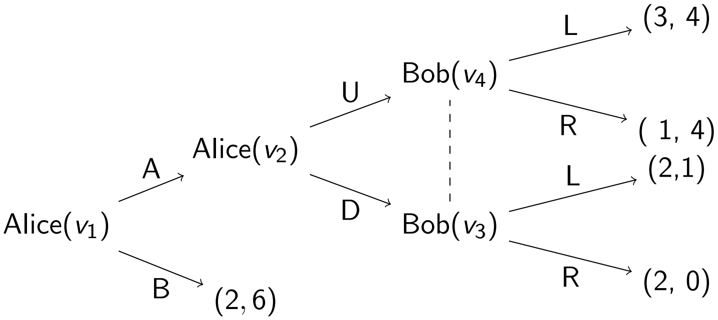

($\{v_3, v_4\}$ is an information set)

This game has two subgames: the game itself, and the game corresponding to the subtree with root v_2. The proper subgame has strategic form:

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

and has a Nash-equilibrium (U, L) resulting in payoff $(3, 4)$. Now at the root of the game Alice would choose to move A because $3 > 2$. The resulting Nash Equilibrium ($v_1 \to A, v_2 \to U$, $v_3, v_4 \to L$) is subgame perfect.

A bargaining process

Alice and Bob will share £1, and Alice starts by issuing a take-it-or-leave-it offer of $0 \leq s \leq 1$. If Bob accepts, Alice gets $1 - s$, Bob gets s, otherwise both get 0. We apply backward induction and see that Bob should accept any offer of $s > 0$. (This is not observed in the real world!)

We change the game, so that if Bob rejects the offer, the game is played again: now £δ is shared and Bob makes a take-it-or-leave it offer of $0 \leq t \leq \delta$. Backward induction now shows Alice should offer δ, and the payoffs are $(1 - \delta, \delta)$.

Assume now that there are three iterations, in which £$1, \delta, \delta^2$ are shared in the 1st, 2nd and third iterations. Alice will offer (almost) 0 in the last iteration, and receive δ^2. Now Bob has to offer Alice δ^2 in the second iteration and he receives $\delta(1 - \delta)$. Now Alice has to offer Bob $\delta(1 - \delta)$ and she receives $1 - \delta(1 - \delta) = 1 - \delta + \delta^2$.

What if this process can go on forever? Now payoffs are $(1/(1 + \delta), \delta/(1 + \delta))$, which for δ close to 1, is close to $(1/2, 1/2)$.

What if Alice and Bob have different discount factors δ_1 and δ_2?

Example: The centipede game

$$\begin{tikzcd}
Alice \arrow[r] \arrow[d] & Bob \arrow[r] \arrow[d] & Alice \arrow[r] \arrow[d] & Bob \arrow[r] \arrow[d] & Alice \arrow[r] \arrow[d] & Bob \arrow[r] \arrow[d] & (256, 64) \\
(4, 1) & (2, 8) & (16, 4) & (8, 32) & (64, 16) & (32, 128) &
\end{tikzcd}$$

The only pure strategy Nash equilibrium is *Down* at every node.

$$\begin{tikzcd}
Alice \arrow[r] \arrow[d] & Bob \arrow[r] \arrow[d] & Alice \arrow[r] \arrow[d] & Bob \arrow[r] \arrow[d] & Alice \arrow[r] \arrow[d] & Bob \arrow[r] \arrow[d] & (11, 11) \\
(10, 0) & (0, 10) & (10, 0) & (0, 10) & (10, 10) & (0, 10) &
\end{tikzcd}$$

Now *Right* at every node is a (tenuous) subgame perfect Nash equilibrium!