SF2972: Game theory

Mark Voorneveld, mark.voorneveld@hhs.se

February 5, 2013

My part of the course treats:

1. Extensive form games.
   This part consists of the first three lectures.
2. Introduction to matching problems.
   This was the topic of the 2012 'Nobel prize in economics'; my fourth lecture.

Jörgen covered parts of Ch. 4 and 5 succinctly; I will do so in more detail. Moreover, I cover Ch. 14. For the fourth lecture, I will refer to some articles. A reading guide will be posted on the course web.

Topic 1: defining games and strategies

Drawing a game tree is usually the most informative way to represent an extensive form game. Here is one with an initial (c)hance move:

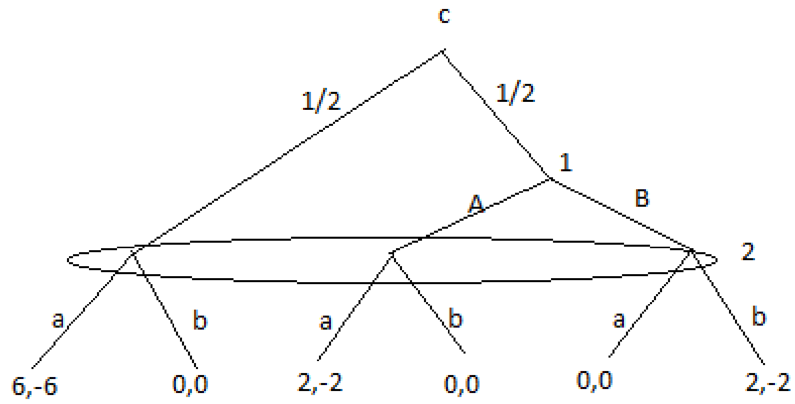

For LaTeX gurus: Is there a neat, quick way to draw game trees with Ti*k*Z?

Extensive form game: formal definition

- A (directed, rooted) tree; i.e. it has a well-defined initial node.
- Nodes can be of three types:
  1. chance nodes: where chance/nature chooses a branch according to a given/known probability distribution;
     Let τ assign to each chance node a prob distr over feasible branches.
  2. decision nodes: where a player chooses a branch;
  3. end nodes: where there are no more decisions to be made and each player i gets a payoff/utility given by a utility function u_i.
- A function P assigns to each decision node a player i in player set N who gets to decide there.
- Decision nodes $P^{-1}(i)$ of player i are partitioned into information sets.
  Nodes in an information set of player i are 'indistinguishable' to player i; this requires, for instance, the same actions in each decision node of the information set.
- If h is an information set of player i, write $P(h) = i$ and let $A(h)$ be the feasible actions in info set h.

Not allowed:

Since nodes in an information set are indistinguishable, information sets like

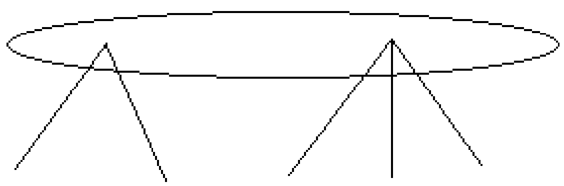

are not allowed: since there are two actions in the left node and three in the right, they are easily distinguishable.

We call an extensive form game *finite* if it has finitely many nodes. An extensive form game has

- *perfect information* if each information set consists of only one node.
- *perfect recall* if each player recalls exactly what he did in the past.
  Formally: on the path from the initial node to a decision node x of player i, list in chronological order which information sets of i were encountered and what i did there. Call this list the experience $X_i(x)$ of i in node x. The game has perfect recall if nodes in the same information set have the same experience.
- otherwise, the game has imperfect information/recall.

Convention: we often characterize nodes in the tree by describing the sequence of actions that leads to them. For instance:

- the initial node of the tree is denoted by $\emptyset$;
- node (a_1, a_2, a_3) is reached after three steps/branches/actions: first a_1, then a_2, then a_3.

Imperfect recall: absentminded driver

Two crossings on your way home. You need to (C)ontinue on the first, (E)xit on the second. But you don't recall *whether* you already passed a crossing.

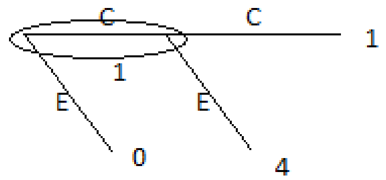

Only one information set, $\{\emptyset, C\}$, but with different experiences:

- in the first node: $X_1(\emptyset) = (\{\emptyset, C\})$
- in the second node:
  $X_1(C) = (\underbrace{\{\emptyset, C\}}_{\text{1's first info set}}, \underbrace{C}_{\text{choice there}}, \underbrace{\{\emptyset, C\}}_{\text{resulting info set}})$
- $X_1(\emptyset) \neq X_1(C)$: imperfect recall!

Second example of imperfect recall

Player 1 forgets the initial choice:

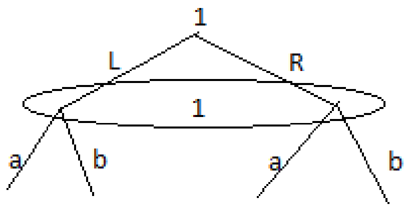

Different experiences in the two nodes of information set $\{L, R\}$:

- in the left node: $X_1(L) = (\underbrace{\emptyset}_{\text{initial node}}, \underbrace{L}_{\text{choice there}}, \underbrace{\{L, R\}}_{\text{resulting info set}})$
- in the right node: $X_1(R) = (\emptyset, R, \{L, R\})$.
- $X_1(L) \neq X_1(R)$: imperfect recall!

Third example of imperfect recall

Player 1 knew the chance move, but forgot it:

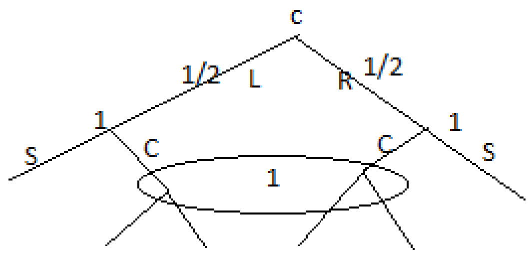

Different experiences in the two nodes of information set $\{(L, C), (R, C)\}$:

- in the left node:
  $X_1((L, C)) = (\underbrace{\{L\}}_{\text{1's first info set}}, \underbrace{C}_{\text{choice there}}, \underbrace{\{(L, C), (R, C)\}}_{\text{resulting info set}})$
- in the right node: $X_1((R, C)) = (\{R\}, C, \{(L, C), (R, C)\})$.
- $X_1((L, C)) \neq X_1((R, C))$: imperfect recall!

Pure, mixed, and behavioral strategies

- A *pure strategy* of player i is a function s_i that assigns to each information set h of player i a feasible action $s_i(h) \in A(h)$.
- A *mixed strategy* of player i is a probability distribution σ_i over i's pure strategies.
  $\sigma_i(s_i) \in [0,1]$ is the prob assigned to pure strategy s_i.
  'Global randomization' at the beginning of the game.
- A *behavioral strategy* of player i is a function b_i that assigns to each information set h of player i a probability distribution over the feasible actions $A(h)$.
  $b_i(h)(a)$ is the prob of action $a \in A(h)$.
  'Local randomization' as play proceeds.

Let us consider the difference between these three kinds of strategies in a few examples.

The difference between mixed and behavioral strategies

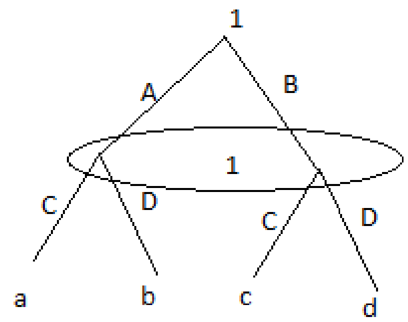

- Imperfect recall; 4 outcomes with payoffs a, b, c, and d.
- Four pure strategies, abbreviated AC, AD, BC, BD.
- Mixed strategies: probability distributions over the 4 pure strategies. A vector $(p_{AC}, p_{AD}, p_{BC}, p_{BD})$ of nonnegative numbers, adding up to one, with p_x the probability assigned to pure strategy $x \in \{AC, AD, BC, BD\}$.

- Behavioral strategies assign to each information set a probability distribution over the available actions. Since pl. 1 has 2 information sets, each with 2 actions, it is summarized by a pair $(p,q) \in [0,1] \times [0,1]$, where $p \in [0,1]$ is the probability assigned to action A in the initial node (and $1-p$ to B) and q is the probability assigned to action C in information set $\{A, B\}$ (and $1-q$ to D).
- Mixed strategy $(1/2, 0, 0, 1/2)$ assigns probability $1/2$ to each of the outcomes a and d. There is no such behavioral strategy:
  - reaching a with positive probability requires that $p, q > 0$;
  - reaching d with positive probability requires $p, q < 1$;
  - hence also b and c are reached with positive probability.

A trickier example: the absentminded driver revisited

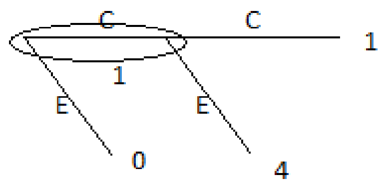

- Pure strategies: C with payoff 1 and E with payoff 0.
- Mixed: let $p \in [0,1]$ be the prob of choosing pure strategy C and $1-p$ the prob of pure strategy E. Expected payoff: p.
- Behavioral: let $q \in [0,1]$ be the prob of choosing action C in the info set and $1-q$ the prob of choosing E in the info set. Expected payoff:

$$0 \cdot (1-q) + 4 \cdot q(1-q) + 1 \cdot q^2 = q(4-3q).$$

- No behavioral strategy is outcome-equivalent with $p = 1/2$ (why?)
- No mixed strategy is outcome-equivalent with $q = 1/2$ (why?)

Outcome-equivalence under perfect recall

Conclude: under imperfect recall, mixed and behavioral strategies might generate different probability distributions over end nodes.
Perfect recall helps to rule this out. We need a few definitions:
Each profile $b = (b_i)_{i \in N}$ of *behavioral strategies* induces an outcome $O(b)$, a probability distribution over end nodes.
How to compute $O(b)$ in finite games?
The probability of reaching end node $x = (a_1, \dots, a_k)$, described by the sequence of actions/branches leading to it, is simply the product of the probabilities of each separate branch:

$$\prod_{\ell=0}^{k-1} b_{P(a_1,\dots,a_\ell)}(a_1, \dots, a_\ell)(a_{\ell+1}).$$

Likewise, each profile $\sigma = (\sigma_i)_{i \in N}$ of *mixed strategies* induces an outcome $O(\sigma)$, a probability distribution over end nodes.
How to compute $O(\sigma)$ in finite games?

- Let $x = (a_1, \dots, a_k)$ be a node, described by the sequence of actions/branches in the game tree leading to it.
- Pure strategy s_i of player i is *consistent with* x if i chooses the actions described by x: for each initial segment $(a_1, \dots, a_\ell)$ with $\ell < k$ and $P(a_1, \dots, a_\ell) = i$:
$$s_i(a_1, \dots, a_\ell) = a_{\ell+1}.$$
- The prob of choosing a pure strategy s_i consistent with x is
$$\pi_i(x) = \sum \sigma_i(s_i),$$
with summation over the s_i consistent with x.
- Similar for nature, whose behavior is given by function τ.
- The probability of reaching end node x is
$$\prod_{i \in N \cup \{c\}} \pi_i(x).$$

A mixed strategy σ_i and a behavioral strategy b_i of player i are *outcome-equivalent* if — given the pure strategies of the remaining players — they give rise to the same outcome:

$$\text{for all } s_{-i}: \quad O(\sigma_i, s_{-i}) = O(b_i, s_{-i}).$$

Theorem (Outcome equivalence under perfect recall)

In a finite extensive form game with perfect recall:

(a) *each behavioral strategy has an outcome-equivalent mixed strategy,*

(b) *each mixed strategy has an outcome-equivalent behavioral strategy.*

Proof sketch:

(a) Given beh. str. b_i, assign to pure strategy s_i the probability
$$\sigma_i(s_i) = \prod_h b_i(h)(s_i(h)),$$
with the product taken over all info sets h of pl i.
Intuition: s_i selects action $s_i(h)$ in information set h. How likely is that?

(b) Given mixed str. σ_i. Consider an info set h of pl i and a feasible action $a \in A(h)$. How should we define $b_i(h)(a)$?
Consider any node x in info set h. The probability of choosing consistent with x is $\pi_i(x)$.
Perfect recall: $\pi_i(x) = \pi_i(y)$ for all $x, y \in h$.
Define
$$b_i(h)(a) = \frac{\pi_i(x, a)}{\pi_i(x)} \quad \text{if } \pi_i(x) > 0 \text{ (and arbitrarily otherwise)}$$
Intuition: conditional on earlier behavior that is consistent with reaching information set h, how likely is i to choose action a?

Example of outcome equivalent strategies

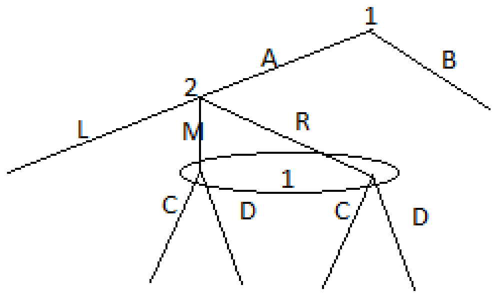

Which behavioral strategy is outcome-equivalent with mixed strategy $(p_{AC}, p_{AD}, p_{BC}, p_{BD})$?
In 1's first information set, the prob that A is chosen is $p_{AC} + p_{AD}$.
In 1's second information set, the prob that C is chosen is computed as the probability of choosing C conditional on earlier behavior that is consistent with this information set being reached:

$$\frac{p_{AC}}{p_{AC} + p_{AD}}.$$

Example of outcome equivalent strategies

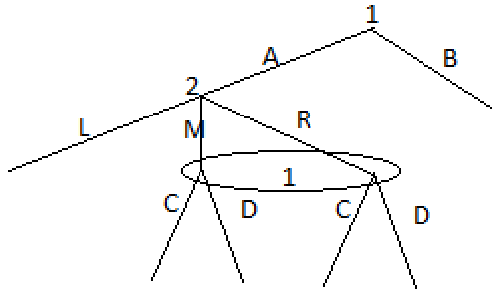

Which mixed strategy is outcome equivalent with the behavioral strategy choosing A with prob p and C with prob q?

$$(p_{AC}, p_{AD}, p_{BC}, p_{BD}) = (pq, p(1-q), (1-p)q, (1-p)(1-q))$$

Homework exercise 1

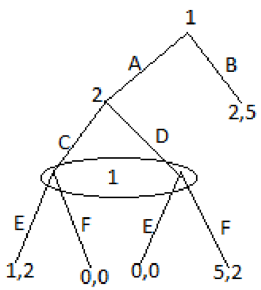

(a) Show that the game above has perfect recall.
(b) For each mixed strategy σ_1 of player 1, find the outcome-equivalent behavioral strategies.
(c) For each behavioral strategy b_1 of player 1, find the outcome-equivalent mixed strategies.

Nash equilibrium

- We can compute, for each profile of pure strategies, the corresponding (expected) payoffs: every extensive form game has a corresponding strategic game. (See already Jörgen's lectures)
- A pure/mixed Nash equilibrium of the extensive form game is then simply a pure/mixed Nash equilibrium of the corresponding strategic game.
- Nash equilibria in behavioral strategies are defined likewise: a profile of behavioral strategies is a Nash equilibrium if no player can achieve a higher expected payoff by unilaterally deviating to a different behavioral strategy.

Theorem (Equilibrium existence)

Every finite extensive form game with perfect recall has a Nash equilibrium in mixed/behavioral strategies.

Proof: finite extensive form game gives finite strategic game, which has a Nash equilibrium in mixed strategies. By outcome-equivalence, we can construct a Nash equilibrium in behavioral strategies.

Strategic form analysis of extensive form games

The extensive form game

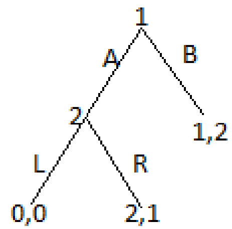

has corresponding strategic form

	L	R
A	$0,0$	$2,1$
B	$1,2$	$1,2$

Pure Nash equilibria: (B, L) and (A, R).
But if pl. 2 is called upon to play, would 2 choose L? This is an implausible choice in the 'subgame' that starts at node A!

Subgame perfect equilibrium

- In an extensive form game with perfect information, let x be a node of the tree that is not an end node. The part of the game tree consisting of all nodes that can be reached from x is called a *subgame*.
- Each game is a subgame of itself. A subgame on a strictly smaller set of nodes is called a *proper subgame*.
- A *subgame perfect equilibrium* is a strategy profile that induces a Nash equilibrium in each subgame.

In the game on the previous slide, only (A, R) is subgame perfect.

Subgame perfect equilibria via backward induction

Subgame perfect equilibria are typically found by backward induction:

1. Start with subgames with only one decision left. Determine the optimal actions there.
2. Next, look at subgames with at most two consecutive decisions left. Conditioning on the previous step, the first player to choose (say i) knows what a 'rational' player will do in the subgame that starts after i's choice, so it is easy to find i's optimal action.
3. Continue with subgames of at most 3 consecutive moves, etc.

Backward induction: example 1

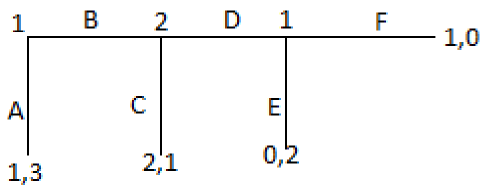

Strategic form:

	C	D
AE	$1,3$	$1,3$
AF	$1,3$	$1,3$
BE	$2,1$	$0,2$
BF	$2,1$	$1,0$

Pure Nash equilibria: (AE, D), (AF, D), and (BF, C).
Subgame perfect equilibrium: (BF, C)

Backward induction: example 2

Dividing 2 indivisible objects. Pl. 1 proposes, pl. 2 accepts or rejects.

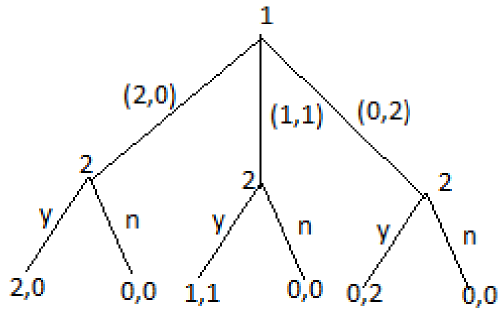

How many pure strategies for player 1? 3
How many pure strategies for player 2? $2^3 = 8$
Subgame perfect equilibria? $((2,0), yyy)$ and $((1,1), nyy)$

Backward induction: the 'rotten kid' game

- A child's action a from some nonempty, finite set A affect both her own payoff $c(a)$ and her parents' payoff $p(a)$; for all $a \in A$ we have $0 \leq c(a) < p(a)$.
- The child is selfish: she cares only about the amount of money she receives.
- Her loving parents care both about how much money they have and how much their child has. Specifically, model the parents as a single player whose utility is the smaller of the amount of money the parents have and the amount the child has. The parents may transfer money to the child (pocket money, trust fund, etc).
- First the child chooses action $a \in A$.
- Then the parents observe the action and decide how much money $x \in [0, p(a)]$ to transfer to the child. The game ends with utility $c(a) + x$ for the child and $\min\{c(a) + x, p(a) - x\}$ to the parents.

Show: in a subgame perfect equilibrium, the child takes an action that maximizes the sum of her private income $c(a)$ and her parents' income $p(a)$. Not so selfish after all!

- In the subgame after action $a \in A$, the parents maximize $\min\{c(a) + x, p(a) - x\}$ over $x \in [0, p(a)]$.
- This is done by choosing x such that $c(a) + x = p(a) - x$, i.e., by $x^*(a) = \frac{1}{2}(p(a) - c(a))$.
- Anticipating this, the child knows that action $a \in A$ leads to transfer $x^*(a)$ and consequently utility $c(a) + x^*(a) = \frac{1}{2}(c(a) + p(a))$. Maximizing this expression is equivalent with maximizing $c(a) + p(a)$.

Finite trees: existence of subgame perfect equilibria

Using backward induction, if there are only finitely many nodes, the first player to move — conditioning on the optimal behavior in the smaller subgames — is optimizing over a finite set: an optimum will always exist. Using this, one can show:

Theorem (Existence of subgame perfect equilibria)

In a finite extensive form game with perfect information, there is always a subgame perfect equilibrium in pure strategies.

That's a pretty nice result:

1. no need to consider randomization
2. no implausible behavior in subgames

As an aside: what if there are infinitely many nodes?

Centipede games

Although subgame perfect equilibria were introduced to rule out implausible behavior in subgames, there are examples where such equilibria lead to outcomes that some people find counterintuitive. This is sometimes corroborated with experimental support. One well-known example consists of Rosenthal's centipede games, characterized by the following properties:

- Players 1 and 2 take turns during at most $2T$ rounds ($T \in \mathbb{N}$).
- At each decision node, the player can choose to (S)top or (C)ontinue.
- The game ends *(i)* if one of the players decides to stop, or *(ii)* if no player has chosen stop after $2T$ periods.
- For each player, the outcome when he stops the game in period t is:
  - *better* than the outcome if the other player stops in period $t+1$ (or the game ends),
  - *worse* than any outcome that is reached if in period $t+1$ the other player continues.

Centipede games

Here is an example of a centipede game with 6 periods:

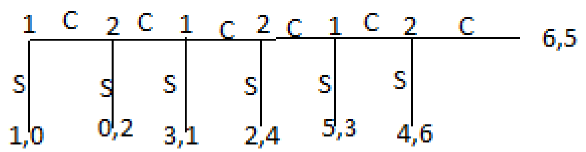

It is tempting to continue the game if you can be sure that the other player does so as well: the longer the game goes on, the higher the payoffs.

But in the unique subgame perfect equilibrium, players choose (S)top in each node. In particular, the game ends immediately in the initial node.

Reason: in the final node, player 2's best reply is to (S)top. Given that 2 (S)tops in the final round, 1's best reply is to stop one period earlier, etc.

There are other Nash equilibria, but they all lead to the same outcome: player 1 ends the game immediately.

- Subgame perfect equilibria in games with perfect information require each player to play a best reply to other players' strategies in each subgame — regardless of whether that subgame is reached or not.
- It is possible to extend the notion of subgame perfect equilibria to games with imperfect information. But the definition of subgames is trickier: information sets must lie entirely outside the subgame or entirely inside the subgame.
- Formally, let x be a (non-end) node and let V^x be the nodes of the tree that can be reached from x. A well-defined subgame starts at x if and only if each information set h of the original game is a subset of V^x or is a subset of its complement.
- Since extensive form games with imperfect information need not have proper subgames, the notion of subgame perfection typically has little 'bite'.

Homework exercise 2

In the game of homework exercise 1:

(a) Find the corresponding strategic game.

(b) Find all pure-strategy Nash equilibria.

(c) What is the outcome of iterated elimination of weakly dominated (pure) strategies?

(d) Find all subgame perfect equilibria (in behavioral strategies).

Another plausible generalization of the notion of subgame perfect equilibria to games with imperfect information would be

require best responses in each information set.

Problem: the best response depends on where in the information set the player believes to be!

The following game has no proper subgames:

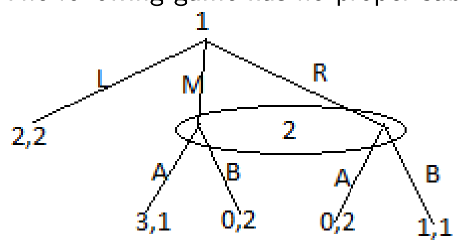

Beliefs affect optimal strategies: consider pl 2 in info set $\{M, R\}$. A is a best response if and only if the player assigns at most prob $1/2$ to being in node M.
Strategies affect reasonable beliefs: If pl 1 assigns to actions (L, M, R) probabilities $(\frac{1}{10}, \frac{3}{10}, \frac{6}{10})$, pl 2 is twice as likely to end up in node R than in node M. Bayes' law requires that he assigns conditional prob $1/3$ to M and $2/3$ to R.
Question: What are reasonable beliefs if 1 chooses L with prob 1?

Assessments

We consider two requirements on beliefs that give different answers to the final question:

1. Bayesian consistency: in information sets that are reached with positive probability, beliefs are determined by Bayes' law. In information sets reached with zero probability, beliefs are allowed to be arbitrary.
2. Consistency: beliefs are determined as a limit of cases where everything happens with positive probability and — consequently — where Bayes' law can be used.

In particular, in both of these notions, we need to define two things: strategies and beliefs over the nodes in the information sets. The difference will lie in the constraints that are imposed.
Formally, consider a finite extensive form game with perfect recall. An *assessment* is a pair (b, β), where

- $b = (b_i)_{i \in N}$ is a profile of behavioral strategies and
- β is a belief system, assigning to each information set h a probability distribution β_h over its nodes.

Two belief requirements

Given node x and behavioral strategies b, let $\mathbb{P}_b(x)$ be the probability that node x is reached using b: it is the product of the probabilities assigned to the branches leading to x. Similarly, if h is an information set, it is reached with probability $\mathbb{P}_b(h) = \sum_{x \in h} \mathbb{P}_b(x)$.
Assessment (b, β) is:

- *Bayesian consistent* if beliefs in information sets reached with positive probability are determined by Bayes' law:

$$\beta_h(x) = \mathbb{P}_b(x) / \mathbb{P}_b(h)$$

  for every info set h with $\mathbb{P}_b(h) > 0$ and every node $x \in h$.
- *consistent* if there is a sequence of Bayesian consistent assessments $(b^m, \beta^m)_{m \in \mathbb{N}}$ with each b^m completely mixed (all actions in all info sets have positive prob) and $\lim_{m \to \infty} (b^m, \beta^m) = (b, \beta)$.

Note: (b, β) consistent $\Rightarrow$ (b, β) Bayesian consistent.

Homework exercise 3

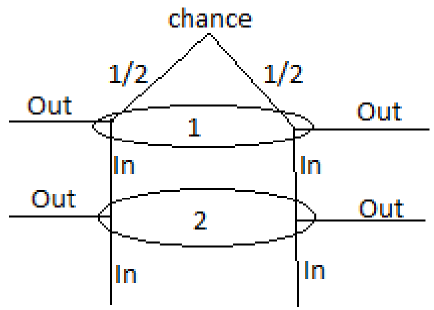

In the game above, where payoffs are omitted since they are irrelevant to the question:

(a) Find all Bayesian consistent assessments (b, β).
(b) Find all consistent assessments (b, β).

Expected payoffs in information sets

Fix assessment (b, β) and an information set h of player i. To formalize the requirement that i plays a best response in info set h, we need to specify i's expected payoff:

1. Conditional on i being in his info set h, belief system β assigns probability $\beta_h(x)$ to being in node $x \in h$.
2. Given such a node x, the probability $\mathbb{P}(e \mid b, x)$ that an end node e is reached, conditional on starting in x and using strategies b is
   - zero if e cannot be reached from x;
   - the product of the probabilities of the corresponding branches from x to e otherwise.
3. In end node e, the payoff to i equals $u_i(e)$.
4. So the expected payoff to agent i in his information set h, given assessment (b, β) is

$$u_i(b_i, b_{-i} \mid h, \beta) = \sum_{x \in h} \beta_h(x) \left(\sum_e \mathbb{P}(e \mid b, x) u_i(e) \right).$$

Sequential rationality

Assessment (b, β) is *sequentially rational* if each player i in each of his information sets h chooses a best response to the belief system β and the strategies of the other players:

$$u_i(b_i, b_{-i} \mid h, \beta) \geq u_i(b_i', b_{-i} \mid h, \beta)$$

for all other behavioral strategies b_i' of player i.
Note:

1. consistency says that beliefs have to make sense given the strategies, without requirements on the strategies;
2. sequential rationality says that strategies have to make sense given the beliefs, without requirements on the beliefs.

Putting the two together, we have:

Perfect Bayesian equilibria and sequential equilibrium

An assessment (b, β) is a *sequential equilibrium* if it is consistent and sequentially rational.

Theorem (Relations between solution concepts for extensive form games)

(a) *Each finite extensive form game with perfect recall has a sequential equilibrium.*

(b) *If assessment (b, β) is a sequential equilibrium, then b is a subgame perfect equilibrium and (hence) a Nash equilibrium.*

Proof sketch

(a) Via perfect equilibria of an auxiliary 'agent-strategic form game':
- Each player i is split up into agents, one agent for each of i's information sets;
- Agents of i have the same preferences as i;
- A mixed strategy in this agent-strategic form game is a behavioral strategy in the original game;
- Consider a completely mixed seq $b^m \to b$ making b a perfect equilibrium
- For each b^m, Bayes' law gives a belief system β^m.
- Drawing a convergent subsequence if necessary, we can show that $\lim_{m\to\infty}(b^m, \beta^m) = (b, \beta)$ is a sequential equilibrium.

(b) Suppose not. Let i have a profitable deviation b_i' in a subgame starting at some node x. Hence, in this subgame there has to be an information set that is reached with positive probability and where i has a profitable deviation, contradicting sequential rationality and correctness of beliefs.

Compute the sequential equilibria of the game below:

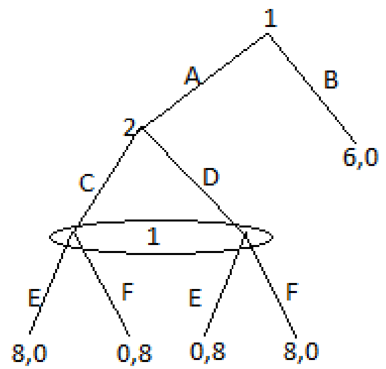

Intuition: What should it be? Player 1 chooses between sure payoffs $(6, 0)$ or the strategic game

	C	D
E	$8, 0$	$0, 8$
F	$0, 8$	$8, 0$

Behavioral strategies $b = (b_1, b_2)$ can be summarized by three probabilities:

1. p, the prob that 1 chooses A in the initial node;
2. q, the prob that 2 chooses C in his information set $\{A\}$;
3. r, the prob that 1 chooses E in information set $\{(A, C), (A, D)\}$.

Belief system β can be summarized by one probability α, the prob assigned to the left node (A, C) in the information set $\{(A, C), (A, D)\}$. Consistency: completely mixed beh. str. have $p, q, r \in (0, 1)$. Bayes' law then gives

$$\alpha = \frac{pq}{pq + p(1-q)} = q,$$

So for each consistent assessment (b, β), it follows that $\alpha = q$. Which of these assessments also satisfies sequential rationality?

Distinguish 3 cases:

1. If $q = 0$, then $\alpha = 0$, so $r = 0$ is 1's unique best reply in the final info set. But if $r = 0$, then $q = 0$ is not a best reply in 2's info set. Contradiction.
2. If $q = 1$, then $\alpha = 1$, so $r = 1$ is 1's unique best reply in the final info set. But if $r = 1$, then $q = 1$ is not a best reply in 2's info set. Contradiction.
3. If $q \in (0, 1)$, rationality in 2's info set $\{A\}$ dictates that both C and D must be optimal. C gives $8(1 - r)$, D gives $8r$, so $r = 1/2$.
   In the info set $\{(A, C), (A, D)\}$ of pl. 1, his expected payoff is
   $$\alpha[8r] + (1 - \alpha)[8(1 - r)] \underbrace{=}_{\alpha = q} = 8 - 8q + 8r(2q - 1).$$
   Choosing $r = 1/2$ is rational only if $q = 1/2$.
   Finally, in the initial node, A gives expected payoff 4 and B gives expected payoff 6, so $p = 0$.

Conclude: there is a unique sequential equilibrium with $p = 0, q = r = \alpha = 1/2$.

Homework exercise 4

Find the sequential equilibria (b, β) of the game in homework exercise 1.

Perfect Bayesian equilibrium

- Economists sometimes use perfect Bayesian equilibria, a notion that is more restrictive than subgame perfection, but less restrictive than sequential equilibrium.
- The intuition is that assessments are derived from strategies following Bayes' law whenever possible, but the exact definition of 'whenever possible' differs.
- Therefore, I will not discuss this notion further: if you have a carefully written game theory paper, the authors will make their equilibrium notion precise.
- In practice, a common requirement is that beliefs have to be Bayesian consistent with strategies in the game itself, but also in its subgames.

Signalling games

1. Chance chooses a type t from some nonempty finite set T according to known prob distr $\mathbb{P}$ with $\mathbb{P}(t) > 0$ for all $t \in T$.
2. Pl. 1 (the sender) observes t and chooses a message $m \in M$ in some nonempty finite set of messages M.
3. Pl. 2 (the receiver) observes m (not t) and chooses an action $a \in A$ in some nonempty finite set of actions A.
4. The game ends with utilities $(u_1(t, m, a), u_2(t, m, a))$.

A pure strategy for player 1 is a function $s_1 : T \to M$ and a pure strategy for player 2 is a function $s_2 : M \to A$.

Separating and pooling equilibria in signalling games

In signalling games, it is common to restrict attention to equilibria equilibria (s_1, s_2, β), where

- s_1 and s_2 are pure strategies;
- assessment (s_1, s_2, β) is Bayesian consistent;
- assessment (s_1, s_2, β) is sequentially rational.

Sometimes it is in the sender's interest to try to communicate her type to the receiver by sending different messages for different types

$$s_1(t) = s_1(t') \quad \text{for all } t, t' \in T.$$

In such cases we call the equilibrium (s_1, s_2, β) a *separating equilibrium*.

In other cases, the sender might want to keep her signal a secret to the receiver and send the same message for each type:

$$s_1(t) \neq s_1(t') \quad \text{for all } t, t' \in T.$$

In such cases we call the equilibrium (s_1, s_2, β) a *pooling equilibrium*.

Signalling games: example

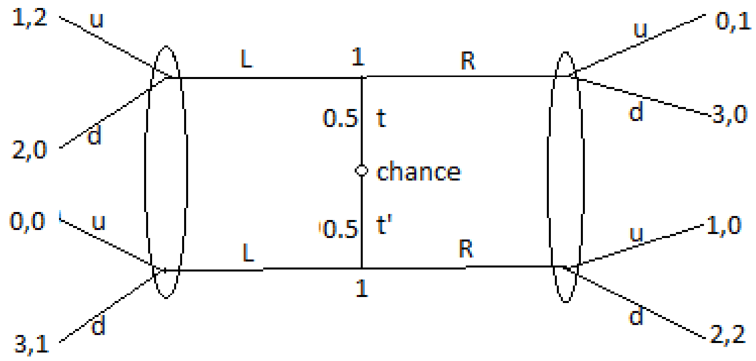

In the signalling game above:

(a) Find the corresponding strategic form game and its pure-strategy Nash equilibria.

(b) Determine (if any) the game's pooling equilibria.

(c) Determine (if any) the game's separating equilibria.

Homework exercise 5

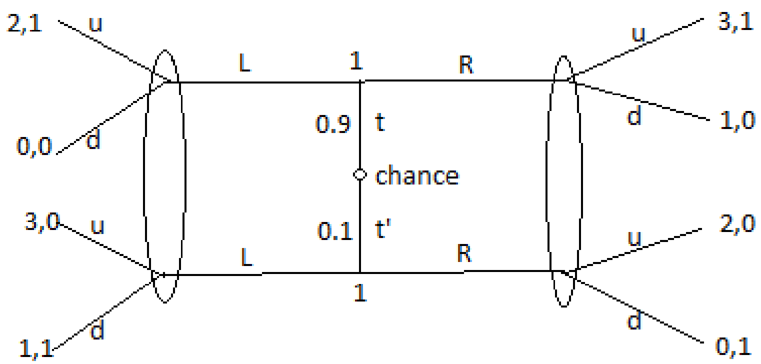

In the signalling game above:

(a) Find the corresponding strategic form game and its pure-strategy Nash equilibria.

(b) Determine (if any) the game's pooling equilibria.

(c) Determine (if any) the game's separating equilibria.

Bayesian games

Bayesian games are special imperfect information games where an initial chance move assigns to each player a privately known type. Knowing their own type, they choose an action (simultaneously, independently) and the game ends. Formally, the timing is as follows:

1. Chance chooses a vector $t = (t_i)_{i \in N}$ of types, one for each player, from a nonempty, finite product set $\times_{i \in N} T_i$ of types, according to known prob distr $\mathbb{P}$ with $\mathbb{P}(t) > 0$ for all $t = (t_i)_{i \in N} \in \times_{i \in N} T_i$.
2. Each player i observes only her own type t_i and chooses an action a_i from some nonempty set A_i.
3. The game ends with utility $u_i(a_1, \dots, a_n, t_1, \dots, t_n)$ to player $i \in N = \{1, \dots, n\}$.

Since $i \in N$ observes only $t_i \in T_i$, a pure strategy of player i is a function $s_i : T_i \to A_i$. Mixed and behavioral strategies are defined likewise.

Bayesian equilibrium

Given her type, i updates her beliefs over other players' types t_{-i} using Bayes' Law: if she is of type t_i^*, she assigns probability

$$\mathbb{P}(t_{-i} \mid t_i^*) = \frac{\mathbb{P}(t_i^*, t_{-i})}{\mathbb{P}\{t \in \times_{j \in N} t_j \mid t_i = t_i^*\}}$$

to the others having types $t_{-i} \in \times_{j \neq i} T_j$. Hence, her expected payoff given type t_i is

$$u_i(s_1, \ldots, s_n \mid t_i) = \sum_{t_{-i} \in T_{-i}} \mathbb{P}(t_{-i} \mid t_i^*) u_i(s_1(t_1), \ldots, s_n(t_n), t_1, \ldots, t_n).$$

It makes sense to require that each player i, for each possible type t_i, chooses her action optimally. That is, $s_i(t_i)$ should solve

$$\max_{a_i} \sum_{t_{-i} \in T_{-i}} \mathbb{P}(t_{-i} \mid t_i) u_i(s_1(t_1), \ldots, a_i, \ldots, s_n(t_n), t_1, \ldots, t_i, \ldots, t_n).$$

Strategies satisfying this requirement form a *Bayesian equilibrium*.