Game Theory

Georg Moser

Institute of Computer Science @ UIBK

Winter 2009

Summary of Last Lecture

Theorem
the axioms (with state neutrality) are satisfied if and only if there exists a (state-independent) utility function u and a conditional-probability function p such that

1. $\max_{x \in X} u(x,t) = 1$ and $\min_{x \in X} u(x,t) = 0$
2. $p(R|T) = p(R|S)p(S|T)$ $\forall$ R, S, T so that $R \subseteq S \subseteq T$ and $S \neq \varnothing$
3. $f \succcurlyeq_S g$ if and only if $E_p(u(f)|S) \geqslant E_p(u(g)|S)$

Content

motivation, introduction to decision theory, decision theory

basic model of game theory, dominated strategies, Bayesian games, Nash equilibrium

two-person zero-sum games, Bayesian equilibrium, sequential equilibria of extensive-form games, computing Nash equilibria, sub-game-perfect equilibria

efficient computation of Nash equilibria, complexity class PPAD, complexity of Nash equilibria, refinements of equilibrium in strategic form, persistent equilibria, games with communication, sender-receiver games

Domination

Question
how to find conditional-probability functions?

Answer
game theory or dominated strategies

Preparation

- a decision-maker has an utility function $u \colon X \times \Omega \to \mathbb{R}$
- re-interpret X as the set of possible decisions
- decision-maker can choose any $x \in X$
- let $p(t) = p(t|\Omega)$, $y \in X$ is good decision if:

$$\sum_{t \in \Omega} p(t) \cdot u(y,t) \geqslant \sum_{t \in \Omega} p(t) \cdot u(x,t) \qquad \forall x \in X$$

Definition

a set of vectors S is convex if for any two vectors p, q also $\lambda p + (1-\lambda)q \in S$; generalise to functions in the standard way

Theorem

given $u\colon X \times \Omega \to \mathbb{R}$ and $y \in X$, then
the set of all probabilities $p \in \Delta(\Omega)$ such that y is optimal is convex

Proof

- suppose y is optimal for decision-maker with beliefs p and q
- $\lambda \in [0,1]$, $r = \lambda p + (1-\lambda)q$

$$\begin{aligned}\sum_{t\in\Omega} r(t)\cdot u(y,t) &\geqslant \lambda \sum_{t\in\Omega} p(t)u(y,t) + (1-\lambda)\sum_{t\in\Omega} q(t)u(y,t) \\ &\geqslant \lambda \sum_{t\in\Omega} p(t)u(x,t) + (1-\lambda)\sum_{t\in\Omega} q(t)u(x,t) \\ &= \sum_{t\in\Omega} r(t)\cdot u(x,t)\end{aligned}$$

■

Example

let $X = \{\alpha, \beta, \gamma\}$, let $\Omega = \{\Theta_1, \Theta_2\}$

- $p(\Theta_1) = 1 - p(\Theta_2)$
-

decision	Θ_1	Θ_2
α	8	1
β	5	3
γ	4	7

- α optimal if

$$\begin{aligned}8p(\Theta_1) + 1(1-p(\Theta_1)) &\geqslant 5p(\Theta_1) + 3(1-p(\Theta_1)) \\ 8p(\Theta_1) + 1(1-p(\Theta_1)) &\geqslant 4p(\Theta_1) + 7(1-p(\Theta_1))\end{aligned}$$

- hence α is optimal iff $p(\Theta_1) \geqslant 0.6$
- similar for γ: $p(\Theta_1) \leqslant 0.6$
- but β is never optimal

Definition — strongly dominated

a decision option that can never be optimal is called strongly dominated

Definition — randomised strategy

- a strategy is any probability distribution over decision options X
- notation: $\sigma = (\sigma(x))_{x\in X}$, $\sigma(x) \in \Delta(X)$

Definition — strongly dominated

an option $y \in X$ is strongly dominated by a randomised strategy σ if

$$\sum_{x\in X} \sigma(x)u(x,t) > u(y,t) \qquad \forall t \in \Omega$$

Theorem

both notions of strongly domination are equivalent

Proof

use optimisation results on linear programming problems, more precisely the duality theorem ■

Games in Extensive Form

Example

- player 1 and 2 put 1€ in a pot
- player 1 draws a card, which is either red or black
- player 1 looks at this card in private and can either raise or pass
- if player 1 passes, then she shows the card
  - if the card is red, then player 1 wins the pot
  - if the card is black, then player 2 wins the pot
- if player 1 raises, she adds another euro
- player 2 can meet or fold
  - if player 2 folds the game ends and player 1 wins the pot
  - if player 2 meets she has to add 1€
- the games continues as above

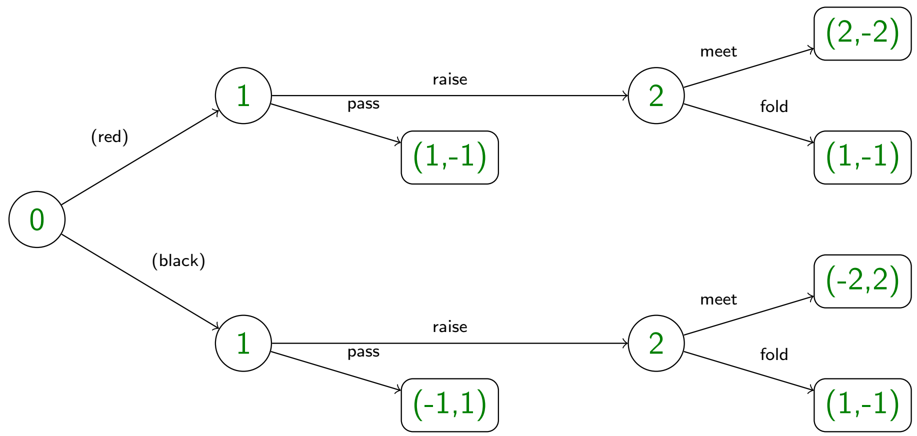

Definition

- node 0 is a chance node
- nodes 1,2 are decision nodes
- the path representing the actual events is called path of play

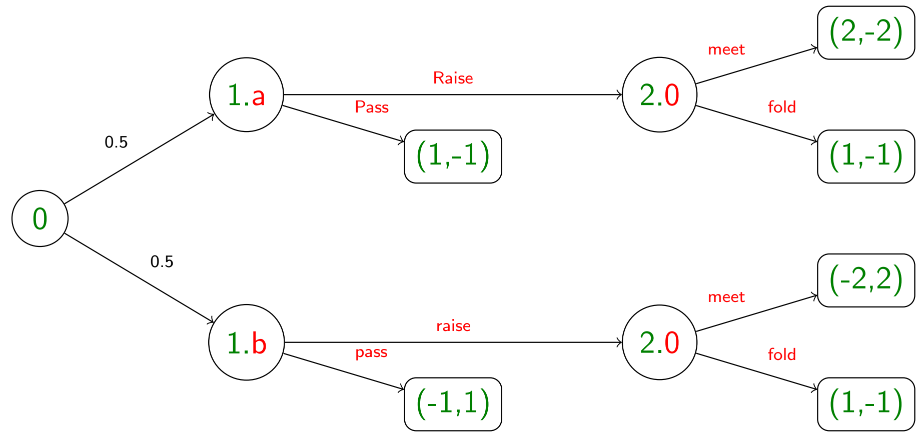

Definition

each decision nodes has two labels

1. the player label
2. the information label

Requirement

the set of move-labels following two nodes must be the same if the two nodes are controlled by the same player in the same information state

n-Person Extensive-Form Game

Definition

an n-person extensive-form game Γ^e is a labelled tree, where also edges are labelled such that

1. each nonterminal node has player label in $\{0, 1, \ldots, n\}$
   nodes labelled with 0 are called chance nodes
   nodes labelled within $\{1, \ldots, n\}$ are called decision nodes
2. edges leaving chance nodes (also called alternatives)
   are labelled with probabilities that sum up to 1
3. player nodes have a second label, the information label
   reflecting the information state
4. each alternative at a player node has a move label
5. each terminal node is labelled with $(u_1, \ldots, u_n)$, the payoff

6. $\forall$ player i,
   $\forall$ nodes x y z controlled by i,
   $\forall$ alternatives b at x
   - suppose y and z have the same information state
     y follows x and b
   - $\exists$ node w, $\exists$ alternative c at w
     such that z follows w and c
   - and w is controlled by player i
     w has the same information label as x
     c has the same move label as b

Question

what does the last condition mean?

Answer

it asserts perfect recall: whenever a player moves, she remembers all the information she knew earlier

No Perfect Recall

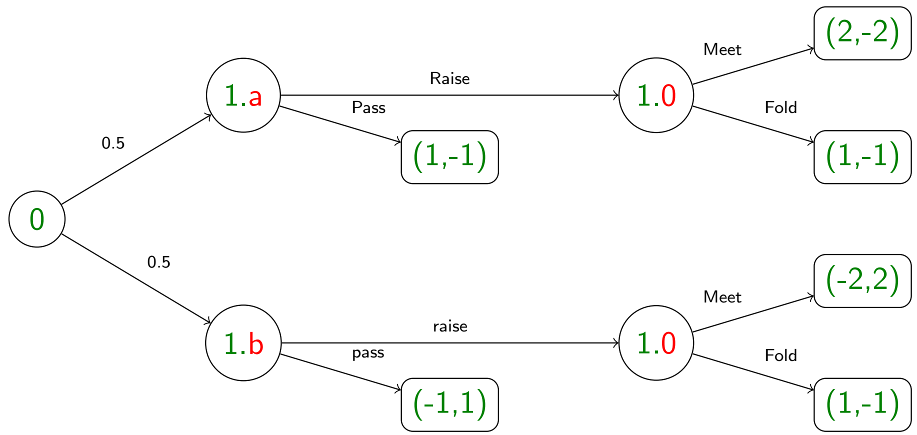

Perfect Information Games

Definition

if no two nodes have the same information state, we say the game has perfect information

Definition

strategy

- S_i is the set of information states per player i
- D_s is the set of possible moves at $s \in S_i$
- the set of strategies for player i is

$$\prod_{s \in S_i} D_s = \underbrace{D_s \times D_s \times \cdot \times D_s}_{S_i\text{-times}}$$

Example

the set of strategies for player 1 can be represented as

$$\{Rr, Rp, Pr, Pp\}$$

Influencing Your Opponent

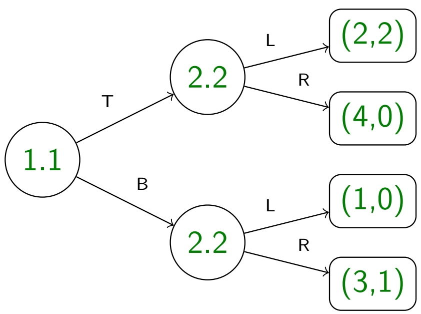

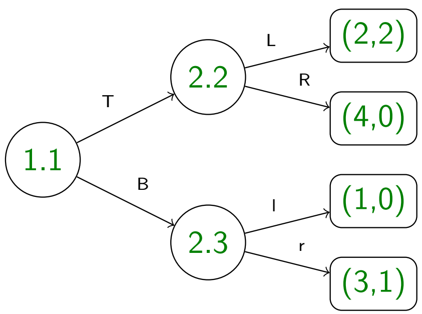

Observation

player 1 profits more, if she chooses T

player 2 doesn't know player 1's choice

Observation

player 1 profits more, if she chooses B

player 2 does know player 1's choice

Example

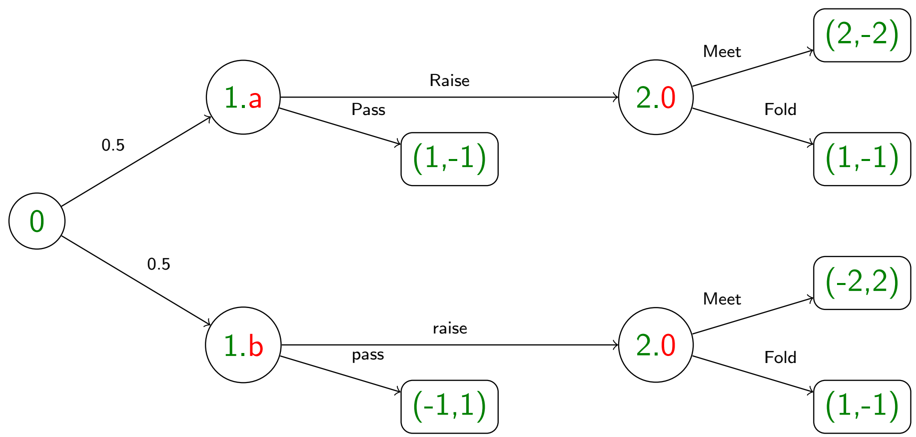

Strategies

$$\underbrace{\{Rr, Rp, Pr, Pp\}}_{\text{for player 1}} \qquad \underbrace{\{M, F\}}_{\text{for player 2}}$$

Strategic-Form Games

Definition

a strategic-form game is a tuple $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ such that

1. N is the set of players
2. for each i: C_i is the set of strategies of player i
3. for each i: $u_i \colon \prod_{i \in N} C_i \to \mathbb{R}$ is the expected utility payoff

a strategic-form game is finite if N and each C_i is finite

Example

consider the card game, suppose player 1 plans to use strategy Rp and player 2 plans to use M

$$u_1(Rp, M) = 2 \cdot \frac{1}{2} + -1 \cdot \frac{1}{2} = 0.5$$

$$u_2(Rp, M) = -2 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2} = -0.5$$

Definition

given a game Γ^e in extensive form, we define the normal representation as strategic-form game $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$:

1. $N = \{1, \ldots, n\}$, if Γ^e is an n-person game
2. for each i: C_i denotes the strategies of each player as defined above
3. we define the expected utility payoff u_i
   - set $C = \prod_{i \in N} C_i$
   - let x be a node in Γ^e
   - let $c \in C$ denote a given strategy profile
   - let $P(x|c)$ denotes the probability that the path of play goes through x, if c is chosen
   - let Ω^* denote the set of all terminal nodes
   - for $x \in \Omega^*$, $w_i(x)$ denotes the payoff for player i
   - set

$$u_i(c) = \sum_{x \in \Omega^*} P(x|c) w_i(x)$$

Example

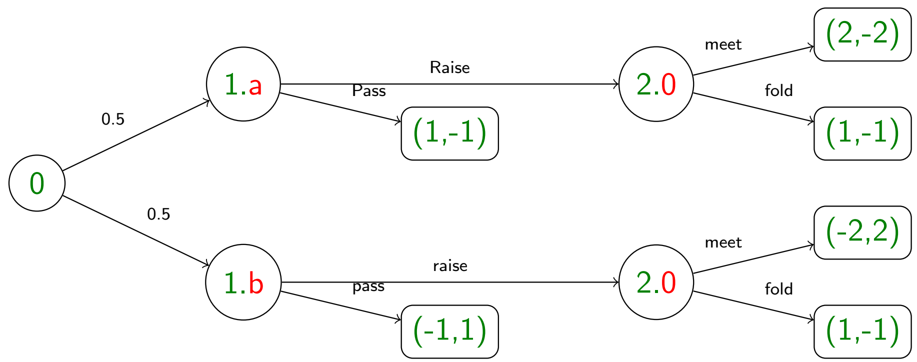

Normal Representation

C_1	C_2: M	F
Rr	$0, 0$	$1, -1$
Rp	$0.5, -0.5$	$0, 0$
Pr	$-0.5, 0.5$	$1, -1$
Pp	$0, 0$	$0, 0$