Lecture notes on game theory

Reading group CSC

March – April 2008

0 Probability Theory

Definition 0.1 (Sample space)**.** *The set of all possible outcomes in an experiment is called the* sample space*.*

Definition 0.2 (Event)**.** *A subset of the sample space is called an* event*.*

Definition 0.3 (Probability for at most countably infinite sets)**.** *Let S be a finite or countably infinite set. A function $p : \mathcal{P}(S) \to \mathbb{R}$ is a probability measure (or a* probability distribution*) if it satisfies the following three axioms:*

1. $\forall E \subset S : 0 \leq p(E) \leq 1$,
2. $p(S) = 1$
3. *For any sequence $E_1, E_2, \ldots \subset S$ with $E_i \cap E_j = \emptyset$ for $i \neq j$,*

$$p\left(\bigcup_{i=1}^{\infty} E_i\right) = \sum_{i=1}^{\infty} p(E_i)$$

We refer to $p(E)$ as the probability *of the event E.*

Exercise 0.4. *Show that $p(\emptyset) = 0$.*

Exercise 0.5. *Show that $p(E^c) = 1 - p(E)$.*

Exercise 0.6. *Show that for $E \subset F$, $p(E) \leq p(F)$.*

Exercise 0.7. *Show that $p(E \cup F) = p(E) + p(F) - p(E \cap F)$.*

Definition 0.8 (Conditional probability)**.** *If $p(F) > 0$ then*

$$p(E|F) = \frac{p(E \cap F)}{p(F)}$$

is called the conditional probability *that E occurs, given that F has occurred.*

Definition 0.9 (Independence)**.** *If $p(E \cap F) = p(E)p(F)$, i.e. $p(E|F) = p(E)$, then we say that E and F are* independent*.*

Definition 0.10 (Random Variable)**.** *A* random variable *is a map from a sample space into a* state space *(typically $\mathbb{R}$) equipped with a probability distribution.*

A random variable is discrete *if it takes on at most countably infinitely many values.*

Definition 0.11 (Expected Value)**.** *If X is a random variable, mapping a sample space S into an at most countably infinite subset of $\mathbb{R}$ equipped with a probability distribution p, then $E(X) = \sum_{x \in X(S)} x \cdot p(x)$ is defined to be the* expected value *of X.*

Exercise 0.12. *The expected value is monotone, in the sense that for $X \leq Y$ $E(X) \leq E(Y)$.*

Exercise 0.13. *Show that the expected value is linear. That is, for random variables X, Y*

- $E(X + c) = E(X) + c$ *for* $c \in \mathbb{R}$.
- $E(X + Y) = E(X) + E(Y)$ *for random variables* X, Y
- $E(aX) = aE(X)$, *for* $a \in \mathbb{R}$.

1 Decision Theoretic Foundations

1.1 Game Theory, Rationality, and Intelligence

Game Theory can be defined as study of mathematical models of conflict and cooperation between intelligent rational decision makers.

A *game* is any social situation involving two or more *players, decision-makers*, or, *individuals.*

The usual assumptions are that players are *rational* and *intelligent*.

A decision-maker is called *rational*, if he makes decisions consistently in pursuit of his own objectives (objective being maximizing expected value of his own payoff).

A player is called *intelligent*, if he knows everything about the game that we know about the game and can make the same inferences about the situation as we can.

1.2 Basic Concepts of Decision Theory

Logical roots of Game Theory are in Bayesian Decision Theory. Game Theory can be considered an extension in that two or more decision-makers are considered.

There are two models, *probability model* and *state-variable model*, which will be combined in a minute. In both models, a decision-maker chooses among *lotteries*. Difference in the models is the definition of a lottery.

	probability model	state-variable model
lottery	probability distrib. over set of prizes	functions from set of possible states into set of prizes
events	objective unknowns	subjective unknowns
examples	gambles depending on coin tossing	stock market, outcome of future sport events

Definition 1.1 (Set of Probability Distributions)**.** *For any finite set Z, the set of probability distributions over Z is denoted by $\Delta(Z)$, defined as follows*

$$\Delta(Z) = \left\{ q : Z \to \mathbb{R} \mid \sum_{y \in Z} q(y) = 1 \text{ and } \forall_{z \in Z} : q(z) \geq 0 \right\}$$

Definition 1.2 (Set of prizes)**.** *Let X denote the set of all* prizes *that the decision-maker could ultimately get. Prizes are mutually exclusive and exhaust the possible consequences of the decision-maker's decisions.*

A prize could be any commodity bundle or resource allocation.

Definition 1.3 (Set of states)**.** *Let Ω denote the set of possible* states*, one of which will be the* true state of the world.

Assumption: X and Ω are finite (to simplify maths.)

Definition 1.4. *A* lottery *is a function*

$$\begin{aligned} f : X \times \Omega &\rightarrow \mathbb{R}^+ \\ (x,t) &\mapsto f(x|t) \end{aligned}$$

such that for fixed $t \in \Omega$, $f(x|t)$ is a probability distribution over X.
We denote the set of all such lotteries by L:

$$L = \{f : \Omega \rightarrow \Delta(X)\}.$$

$f(x|t)$ is to be interpreted as the objective conditional probability of getting prize x in lottery f given that t is the true state of the world. For this to make sense, the state must be defined broadly enough to summarize all subjective unknowns that might influence the prize to be received.

Thus the definition of lottery allows the representation of any gamble in which the prize may depend on both objective and subjective unknowns.

Definition 1.5 (World events)**.** *The information that a decision maker has of the world is described by events, subset of Ω. The set of all events is denoted by Ξ.*

$$\Xi = \{S \,|\, S \subset \Omega, S \neq \emptyset\}$$

We make the assumption that a decision maker has well-defined preferences over lotteries given any event in Ξ.

Definition 1.6 (Preferences in lotteries)**.** *Let $f, g \in L$, $S \in \Xi$. We write*

- *$f \gtrsim_S g$, iff f is at least as desirable as g for the decision-maker if he learned that the true state of the world is in the set S.*
- *$f \sim_S g$, iff $f \gtrsim_S g$ and $g \gtrsim_S f$.*
- *$f >_S g$, iff $f \gtrsim_S g$ and $g \not\gtrsim_S f$.*

We leave out the subscript S when $S = \Omega$.

Definition 1.7. *For $f, g \in L$, $\alpha \in [0,1]$, we write $\alpha f + (1-\alpha)g$ for the lottery $\alpha f(x|t) + (1-\alpha)g(x|t)$*

Example 1.8. *Let an urn be filled with N balls αN of which are white balls and $(1-\alpha)N$ black balls. Suppose the decision maker gets to play lottery f if a white ball gets drawn from the urn, and lottery g if a black ball gets drawn. Thus the decision-maker's ultimate probability to win prize x, if t is the true state of the world, is $\alpha f(x|t) + (1-\alpha)g(x|t)$. So $\alpha f + (1-\alpha)g$ represents the compound lottery built up from f and g by this random lottery selection process.*

Definition 1.9. *For any prize $x \in X$, we write $[x]$ for the lottery that always gives prize x for sure.*

$$\forall t \in \Omega : [x](y|t) = \begin{cases} 1 & y = x \\ 0 & y \neq x \end{cases}$$

Example 1.10. *$\alpha[x] + (1-\alpha)[y]$ is the lottery which gives prize x with probability α and prize y with probability $1-\alpha$.*

1.3 Axioms

Note that the following set of axioms is not minimal. It's just a collection of axioms from literature.

The following axioms represent a rational decision-maker's preferences. In the following the quantifiers are $\forall e, f, g, h \in L$, $\forall S, T \in \Xi$, $\forall 0 \leq \alpha, \beta \leq 1$.

Preferences form a complete transitive order over the set of lotteries:

Axiom 1.11 (Completeness/Transitivity)**.**

$$f \gtrsim_S g \text{ or } g \gtrsim_S f$$

$$f \gtrsim_S g \wedge g \gtrsim_S h \implies f \gtrsim_S h$$

Exercise 1.12. *Transitivity of* $\sim_S$

Exercise 1.13.

$$f >_S g \wedge g \gtrsim_S h \implies f >_S h$$

Only possible states are relevant to the decision-maker.

Axiom 1.14 (Relevance)**.** *If* $f(\cdot|t) = g(\cdot|t)$ *for all* $t \in S$*, then* $f \sim_S g$*.*

A higher probability of getting a better lottery is better.

Axiom 1.15 (Monotonicity)**.** *If* $f >_S g$, $0 \leq \beta < \alpha \leq 1$*, then*

$$\alpha f + (1-\alpha)g >_S \beta f + (1-\beta)g$$

Any lottery whose preference is nested between two other lotteries can be written as combination of them.

Axiom 1.16 (Continuity)**.** *If* $f \gtrsim_S g \gtrsim_S h$*, then there is* $0 \leq \gamma \leq 1$ *such that* $g \sim_S \gamma f + (1-\gamma)h$

The following axioms are also known as independence or sure-thing axioms. They are quite important in that they restrict the decision-maker's preference system heavier than the others. The axioms express that preference of alternatives should be consistent independent of additional information.

In other words, if we prefer option 1 over option 2 when an event occurs, but also when the event does not occur, then we'll prefer option 1 even if we don't know whether the event occurred.

Axiom 1.17 ((Strict) Objective Substitution)**.** *If* $e \gtrsim_S f$ *and* $g \gtrsim_S h$, $0 \leq \alpha \leq 1$*, then*

$$\alpha e + (1-\alpha)g \gtrsim_S \alpha f + (1-\alpha)h.$$

Strict version: If $e >_S f$ *and* $g \gtrsim_S h$, $0 < \alpha \leq 1$*, then*

$$\alpha e + (1-\alpha)g >_S \alpha f + (1-\alpha)h.$$

Axiom 1.18 ((Strict) Subjective Substitution)**.** *If* $f \gtrsim_S g$ *and* $f \gtrsim_T g$, $S \cap T = \emptyset$*, then*

$$f \gtrsim_{S \cup T} g.$$

Strict version: If $f <_S g$ *and* $f <_T g$, $S \cap T = \emptyset$*, then*

$$f <_{S \cup T} g.$$

Example 1.19 (Trouble without Substitution)**.** *Given:* $[x] > [y]$, $.5[y] + .5[z] > [w] > .5[x] + .5[z]$ *(violating substitution with* x, y*).*

Consider following situation:

1. *Decide whether to take* w *or not.*
2. *If* w *is not taken, coin toss: heads gives* z
3. *If tails, decide between* x *and* y*.*

There are three possible strategies:

1. *Take* w*, leading to* $[w]$*.*

2. Take x if tails, leading to $.5[x] + .5[z]$.

3. Take y if tails, leading to $.5[y] + .5[z]$.

According to ranking above, a rational decision maker likes third option best. But if he starts following that strategy, when tails comes up, he should prefer x over y, indicating a preference for second strategy. But the second strategy is worse than w, so maybe he should take w!

Without the last two axioms, one would need to discuss whether strategies can be changed after committing to them (leading to the third strategy, at first). If not, we would need to discuss whether future inconsistencies can be foreseen (leading to choice of $[w]$) or not (leading to $.5[x] + .5[z]$)

The next axiom asserts that the decision-maker is never indifferent between all prices. (Avoiding states where nothing of interest would happen, a regularity condition.)

Axiom 1.20 (Interest/Regularity)**.** *For every state $t \in \Omega$ there exist prizes $x, y \in X$ such that $[x] >_{\{t\}} [y]$.*

Finally an optional axiom, which will not be used for the main theorem. It states that the same preference ordering will be maintained across all states.

Axiom 1.21 (State neutrality)**.** *For all $r, t \in \Omega$, if $f(\cdot|r) = f(\cdot|t)$ and $g(\cdot|r) = g(\cdot|t)$ and $f \gtrsim_{\{r\}} g$, then $f \gtrsim_{\{t\}} g$.*

1.4 The Expected-Utility Maximization Theorem

Definition 1.22 (Conditional probability)**.** *A conditional-probability function on Ω is a function $p : \Xi \to \Delta(\Omega)$ such that $p(t|S) = 0$, if $t \notin S$, and $\sum_{r \in S} p(r|S) = 1$. ($S \in \Xi$). We write $p(R|S) = \sum_{r \in R} p(r|S)$ for $R \subset \Omega$.*

Definition 1.23 (Utility function)**.** *A* utility function *is any function $u : X \times \Omega \to \mathbb{R}$. u is* state independent *if there is $U : X \to \mathbb{R}$ such that $u(x, t) = U(x)$.*

Definition 1.24 (Expected Utility Value)**.** *For any conditional-probability function p, utility function u, lottery f, event $S \in \Xi$, the expected utility value of the prize determined by f when $p(\cdot|S)$ is the probability distribution for the true state of the world is*

$$E_p(u(f)|S) = \sum_{t \in S} p(t|S) \sum_{x \in X} u(x, t) f(x|t).$$

The following theorem asserts that for any rational decision-maker there is a way of assigning utility numbers to the possible outcomes such that he will always choose the option that maximizes his expected utility.

Theorem 1.25. *All but the last axiom are jointly satisfied if and only if there exists a utility function $u : X \times \Omega \to \mathbb{R}$ and a conditional-probability function $p : \Xi \to \Delta(\Omega)$ such that*

$$\forall t \in \Omega : \max_{x \in X} u(x, t) = 1 \text{ and } \min_{x \in X} u(x, t) = 0 \tag{1}$$

$$\forall R \subset S \subset T \subset \Omega, S \neq \emptyset : p(R|T) = p(R|S) p(S|T) \tag{2}$$

$$\forall f, g \in L, S \in \Xi : f \gtrsim_S g \Leftrightarrow E_p(u(f)|S) \geq E_p(u(g)|S) \tag{3}$$

In addition, the last axiom is satisfied if and only if (1)-(3) can be satisfied with a state-independent utility function.

(1) is just a normalization, (2) is Bayes' formula. (3) states that the decision-maker always prefers lotteries with a higher expected utility. Thus, once we have assessed u and p, we can predict the decision-maker's optimal choice in any decision making situation.

We now consider a procedure [Raiffa 1968] for assessing $u(x, t)$ and $p(t|S)$.

a_1: lottery that gives the decision-maker best prizes in every state.
a_0: lottery that gives the decision-maker worst prizes in every state.
Thus: $\forall t \in \Omega \exists y, z \in X : a_1(y|t) = 1 = a_0(z|t)$ and $\forall x \in X : y \gtrsim_{\{t\}} x \gtrsim_{\{t\}} z$.
Such best and worst prizes exist in every state by finiteness of X and transitive ordering of $\gtrsim$.

We define some special lotteries, b_S and $c_{x,t}$ (c only needed in proof of theorem) with the following properties: For $S \in \Xi$

$b_S(\cdot|t) = \begin{cases} a_1(\cdot|t) & t \in S \\ a_0(\cdot|t) & t \notin S \end{cases}$ gives the best possible prize when $t \in S$, and the worst possible else.

$c_{x,t}(\cdot|r) = \begin{cases} [x](\cdot|r) & r = t \\ a_0(\cdot|r) & r \neq t \end{cases}$ gives the worst prize, except when in state t.

The procedure to asses p and u is now as follows. Ask decision-maker:

- for each x and t, for what β is decision-maker indifferent between $[x]$ and $\beta a_1 + (1-\beta)a_0$, if t was true state of the world. (Number exists by continuity axiom).

  This defines $u(x,t)$ such that

  $$[x] \sim_{\{t\}} u(x,t)a_1 + (1 - u(x,t))a_0$$

- for each t and S, for what γ is decision-maker indifferent between $b_{\{t\}}$ and $\gamma a_1 + (1-\gamma)a_0$, if the true state of the world was in S. (Number exists by continuity axiom.)

  This defines $p(t|S)$ such that

  $$b_{\{t\}} \sim_S p(t|S)a_1 + (1 - p(t|S))a_0$$

  (Subjective substitution guarantees that $a_1 \gtrsim_S b_{\{t\}} \gtrsim_S a_0$.)

This is a finite procedure, because X and Ω are finite.

1.5 Equivalent Representations

Definition 1.26 (Preference representation). *Given $S \in \Xi$, we say that a utility function u and a conditional-probability function q* represent *the preference ordering $\gtrsim_S$, iff*

$$\forall f, g \in L : E_q(u(f)|S) \geq E_q(u(g)|S) \Leftrightarrow f \gtrsim_S g$$

Consider (1). Without it, we can have several utility functions with equivalent representations. Any economic or other theory, should not change based on different utility functions with equivalent representations. Thus want to be able to recognize equivalent representations.

Theorem 1.27. *Let notation be as in the preceding theorem. Then v, q represent the preference ordering $\gtrsim_S$ if and only if there exists $A \in \mathbb{R}^+$ and $B : S \to \mathbb{R}$ such that for all $t \in S$, $x \in X$*

$$q(t|S)v(x,t) = Ap(t|S)u(x,t) + B(t)$$

Theorem 1.28. *Let notation be as in the preceding theorem, but u, v are additionally assumed to be state-independent, and decision-maker's preference assumed to satisfy the last axiom above. Then for all $t \in S$, $x \in X$,*

$$q(t|S) = p(t|S)$$

and there are $A, C \in \mathbb{R}$, $A > 0$ such that

$$v(x) = Au(x) + C$$

1.6 Bayesian Conditional-Probability Systems

Definition 1.29 ((Bayesian) Conditional-probability System)**.** *A conditional probability system on the finite set* Ω *is any conditional-probability function* p *on* Ω *which satisfies* Bayes' formula*:* $\forall S \in \Omega : p(\cdot|S) \in \Delta(\Omega)$ *s.t.* $p(S|S) = 1$ *and* $p(R|T) = p(R|S)p(S|T)$ *for all* $R \subset S$, $T \supset S$.

$\Delta^*(\Omega)$ *denotes the set of all Bayesian conditional-probability systems on* Ω.

Definition 1.30. $\Delta^0(Z) = \{q \in \Delta(Z) | q(z) > 0, \forall z \in Z\}$

Every $\hat{p} \in \Delta^0(\Omega)$ gives rise to $p \in \Delta^*(\Omega)$ via

$$p(t|S) = \begin{cases} \dfrac{\hat{p}(t)}{\sum\limits_{r\in S} \hat{p}(r)} & t \in S \\ \\ 0 & t \notin S \end{cases}$$

Conditional-probability systems generated this way do not cover all of $\Delta^*(\Omega)$, but $\Delta^*(\Omega)$ can be obtained as the limit

$$p(t|S) = \begin{cases} \lim\limits_{k\to\infty} \dfrac{\hat{p}^k}{\sum\limits_{r\in S} \hat{p}^k(r)} & t \in S \\ \\ 0 & t \notin S \end{cases}$$

of conditional-probability distributions $\{\hat{p}^k\}_{k=1}^{\infty}$ in $\Delta^0(\Omega)$ obtained this way.

1.7 Limitations of the Bayesian Model

Three examples:

Example 1.31 (Challenging objective substitution)**.** *Let* $X = \{\$12, \$1, \$0\}$, $f_1 = .1[\$12] + .9[\$0]$, $f_2 = .11[\$1] + .89[\$0]$, $f_3 = [\$1]$, $f_4 = .1[\$12] + .89[\$1] + .01[\$0]$.

Many people's preference is given by $f_1 > f_2$, $f_3 > f_4$. *But that's not satisfying strict objective substitution, as* $.5f_1 + .5f_3 = .5f_2 + .5f_4$.

Example 1.32 (Challenging subjective probability)**.** *Let* $X = \{-\$100, \$100\}$, $\Omega = \{A, N\}$.

Let $b_A(\$100|A) = 1 = b_A(-\$100|N)$ *be a bet that a team from the American league wins the next All-star game.*

Let $b_N(\$100|N) = 1 = b_N(-\$100|A)$ *be a bet that a team from the National league wins the next All-star game.*

Most people who don't know much about baseball, will have the following preference: $.5[\$100] + .5[-\$100] > b_A$ *and* $.5[\$100] + .5[-\$100] > b_N$, *i.e. would prefer betting on a fair coin toss.*

But this does not work for any subjective probability distribution over Ω, *since at least one state must have a probability greater or equal* $.5$.

This can also be seen as a violation of the strict objective substitution: $.5b_A + .5b_N = .5[\$100] + .5[-\$100]$ *Thus with* $.5b_A + .5b_N > b_A$ *and* $.5b_A + .5b_N > b_N$ *we obtain* $.5b_A + .5b_N > .5b_A + .5b_N$.

Example 1.33 (Difficulty of constructing a model of decision-making)**.**

Theater scenario A *You lost your* \$40 *ticket. Do you buy another* \$40 *ticket?*

Theater scenario B *You have no ticket to begin with, but lost* \$40 *cash on the way to the theater. Do you still buy a* \$40 *ticket (with your cc)?*

Even if both scenarios set you back \$80, *people will behave differently in the two scenarios.*

Some of these problems can be avoided with *salient perturbations*, i.e. small changes to the problem which place the decision-maker into a more familiar scenario for example. Such a perturbation could be a tip from a friend about the Baseball leagues in the second-to-last example.

2 Basic models

We want to represent all possible ways in which a game can be played. There are two forms of representing games:

- the *extensive* form (Kuhn, 1953)
- the *strategic* (or normal) form (to be generalized to the *Bayesian* form).

2.1 Extensive form

Example 2.1. *Consider the following simple card game: Assume a shuffled deck of red and black cards. Two players (A and B) each put a dollar in the pot.*

A draws a card and has two options:

- ***fold**, then A shows his card;*
  - *card is red → A takes pot*
  - *card is black → B takes pot*
- ***raise**, then A adds a dollar to the pot and B has two options:*
  - ***pass** → A takes pot*
  - ***meet** then B adds a dollar to the pot and A shows his card;*
    - *card is red → A takes pot*
    - *card is black → B takes pot*

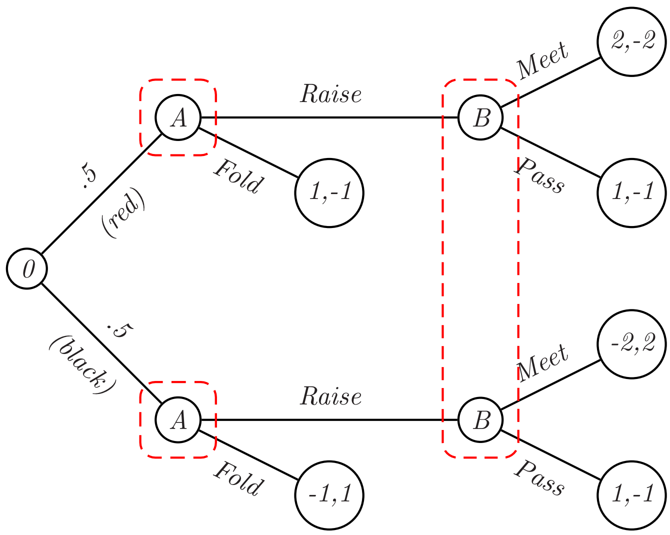

A game representation has three types of nodes:

- decision nodes (in which a player takes a decision to perform a move),
- choice nodes (in which the next state is randomly determined according to some probability distribution), and
- final nodes (in which the pay-out for each of the players is determined).

In order to express that a player cannot make a distinction between certain nodes, we group decision nodes together (such a group is called an *information state*).

Definition 2.2. *Let N be the set of players, of size n. Let U be the set of possible utilities. Let $Prob = [0..1]$ denote the set of probabilities. Let S denote the set of information states. Let M denote the set of moves. An* ***n-person extensive-form game*** *is a rooted node-labeled, edge-labeled tree $\Gamma^e = (V_d, l_d, st, V_c, V_f, l_f, r, E_d, mv, E_c, ch)$, such that (setting $V = V_d \cup V_c \cup V_f$):*

- *V_d is a set of decision nodes, controlled by a player, $l_d : V_d \to N$;*
- *$st : V_d \to S$ determines the information state of the player controlling the node;*
- *V_c is a set of chance nodes;*
- *V_f is a set of final nodes, $l_f : V_f \to (N \to U)$ determines the payoff for each of the players;*
- *$r \in V$ is the root of the tree;*
- *$E_d \subseteq V_d \times V$ is the set of move edges, $mv : E_d \to M$ determines the move, if $mv((v, v')) = mv((v, v''))$ then $v' = v''$;*
- *$E_c \subseteq V_c \times V$ is the set of chance edges;*
- *$ch : E_c \to Prob$ determines the probability of a chance edge, s.t. for all $v \in V_c$, $\Sigma\{ch((v, v')) \mid v' \in V \wedge (v, v') \in E_c\} = 1$*
- *For all $v, v' \in V_d$, if $l_d(v) = l_d(v')$ and $st(v) = st(v')$ then $moves(v) = moves(v')$, where $moves(v) = \{mv((v, v')) \mid v' \in V \wedge (v, v') \in E_d\}$;*

Furthermore, we assume *perfect recall*. A game is said to have perfect recall if each player knows and remembers everything he or she did (but not necessarily what the opponent did). In the model this means that if a player cannot make a distinction between two states, then the history of steps leading to these two states must also be indistinguishable to him.

Let $path(v)$ denote the sequence of edges from the root to node v.

Definition 2.3. *A game Γ^e has* perfect recall *if $\forall x, y \in V_d \quad l_d(x) = l_d(y) \wedge st(x) = st(y) \to path(x) \sim_{l_d(x)} path(y)$, where $\sim_i$ $(i \in N)$ is inductively defined by*

- *$\lambda \sim_i \lambda$ (where λ is the empty sequence of edges).*
- *$p \cdot (v, v') \not\sim_i \lambda$ and $\lambda \not\sim_i p \cdot (v, v')$ (where p a sequence of edges).*
- *$p \cdot (v, v') \sim_i q \cdot (w, w') = p \sim_i q \wedge (ch(v, v') > 0 \leftrightarrow ch(w, w') > 0)$ if $v \in V_c$.*
- *$p \cdot (v, v') \sim_i q \cdot (w, w') = p \sim_i q \wedge mv(v, v') = mv(w, w') \wedge l_d(v) = l_d(w) \wedge st(v) = st(w)$ if $v \in V_d$ and $i = l_d(v) \vee i = l_d(w)$.*
- *$p \cdot (v, v') \sim_i q \cdot (w, w') = p \sim_i q \wedge mv(v, v') = mv(w, w') \wedge l_d(v) = l_d(w)$ if $v \in V_d$ and $i \neq l_d(v) \wedge i \neq l_d(w)$.*

In the above definition we assumed that a player can observe chance events, so he can distinguish between one chance event and a sequence of two chance events. However, he can only observe the outcome of the chance event, not the probability with which it happened. Further, we assumed that a player can observe which players make which moves. Reducing the observational power of a player will lead to slightly different definitions of perfect recall.

If no two nodes have the same information state, then we say that the game has *perfect information.*

A *strategy* is any rule for determining a move at every possible information state in the game.

Definition 2.4. *If S_i denotes the set of information states player i can be in, then $Strategy \in \Pi_{s \in S_i} moves(s)$. (Where $moves(s)$ is the set of moves possible from information state s.)*

Example 2.5. *In the running example, player A has two information states (* red *and* black*) and in each state he has two moves (*raise *and* fold*). So, the four possible strategies are (*raise,raise*), (*raise,fold*), (*fold,raise*), (*fold,fold*). For instance, (*raise,fold*) denotes the strategy in which player A raises if the card is red and folds if the card is black. Because player B has only one information state, he has two strategies:* meet *and* pass.

Exercise 2.6. *Which strategies are possible in the following two games? Which strategies would be smartest?*

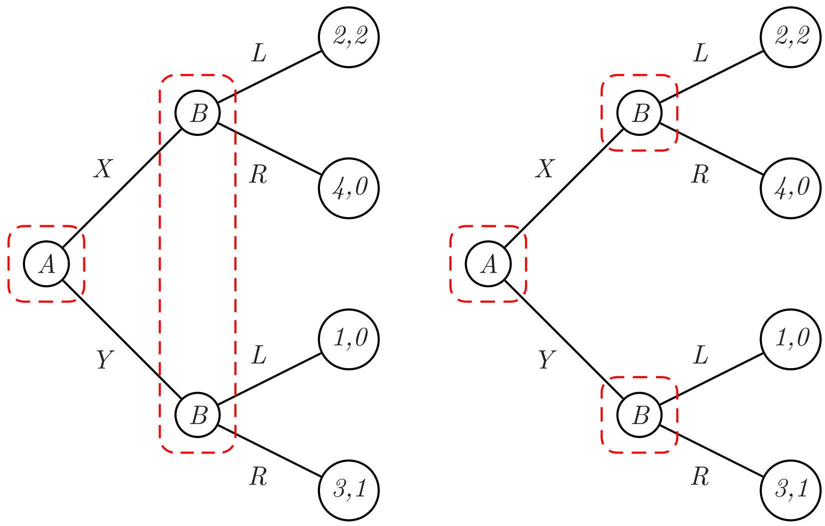

2.2 Strategic form

The extensive representation is a dynamic model of a game, in the sense that it includes a full description of the sequences of moves. In a *strategic* representation we abstract from the dynamic behaviour and assume that all players choose their strategies simultaneously before the game starts. The payoff is then only dependent on the set of chosen strategies.

Definition 2.7. *A* strategic-form *game is a tuple* $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$ *such that*

- N *is a nonempty set of players;*
- $(C_i)_{i\in N}$ *is a set of non-empty strategy sets for every player (we define* $C = \Pi_{i\in N} C_i$*, as the set of strategy profiles);*
- $(u_i)_{i\in N}$ *is a set of utility functions for every player, with* $u_i : C \to \mathbb{R}$*.*

Example 2.8. *Recall that the four possible strategies for player A in the running example are (*raise,raise*), (*raise,fold*), (*fold,raise*), (*fold,fold*) and that the two strategies for player B are* meet *and* pass.

Given that the chance of drawing a red card is equal to the chance of drawing a black card, we can calculate the expected payoff for each combination of strategies, e.g.

$$\begin{aligned} u_A(RF, M) &= 2 \times 0.5 + (-1) \times 0.5 = 0.5 \\ u_B(RF, M) &= -2 \times 0.5 + 1 \times 0.5 = -0.5 \end{aligned}$$

The complete utility function can be displayed in a table. This is the standard notation for the strategic (or normal) form of a game.

A	*B* :	*M*	*P*
RR		*0,0*	*1,-1*
RF		*0.5,-0.5*	*0,0*
FR		*-0.5,0.5*	*1,-1*
FF		*0,0*	*0,0*

Algorithm 2.9. *Every game in extensive form* Γ^e *can be transformed into a game in strategic form* Γ.

- *Determine all strategies of* Γ^e*, this is the set* $(C_i)_{i\in N}$ *of strategies of* Γ.
- *For every strategy profile* $c \in C$ *we attribute every node* v *in* Γ *with the probability* $P(v \mid c)$*, which is the probability of the path to* v*, calculated inductively as follows:*
  - *if* $v = r$ *then* $P(v \mid c) = 1$;
  - *if* $(v', v) \in E_c$ *then* $P(v \mid c) = ch(v', v) \cdot P(v' \mid c)$;
  - *if* $(v', v) \in E_d$ *and* $l_d(v') = i$, $st(v') = \sigma$, *then*
    - *if* $c_i(\sigma) = mv(v', v)$ *then* $P(v \mid c) = P(v' \mid c)$;
    - *if* $c_i(\sigma) \neq mv(v', v)$ *then* $P(v \mid c) = 0$;
- *Determine the utility of each strategy by* $u_i(c) = \Sigma_{v\in V_f} P(v \mid c) \cdot l_f(v)(i)$

Exercise 2.10. *Determine the strategic forms of the games in exercise 2.6. (See p. 49/50, Table 2.2/2.3.)*

2.3 Equivalence of strategic form games

In line with the definition of equivalence on utility functions (see Section 1.5), we define equivalence of games.

Definition 2.11. *Two games in strategic form* $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$ *and* $\Gamma' = (N, (C_i)_{i\in N}, (u'_i)_{i\in N})$ *are* fully equivalent *iff* $\forall_{i\in N} \exists_{A>0,B} \forall_{c\in C} \quad u'_i(c) = Au_i(c) + B$

Corollary 2.12. *(of Theorem 1.27) Let notation be as in the preceding definition. Let* $\Delta(C)$ *denote the set of probability distributions over the set of strategy profiles.* Γ *and* Γ' *are fully equivalent iff*

$$\forall_{i\in N} \forall_{\mu,\lambda\in\Delta(C)} \quad \Sigma_{c\in C}\mu(c)u_i(c) \geq \Sigma_{c\in C}\lambda(c)u_i(c) \quad \leftrightarrow \quad \Sigma_{c\in C}\mu(c)u'_i(c) \geq \Sigma_{c\in C}\lambda(c)u'_i(c)$$

(in words: Γ *and* Γ' *are fully equivalent if each player would prefer probability distribution* μ *over* λ *in the game* Γ *if and only if he would prefer* μ *over* λ *in the game* Γ'*.)*

An alternative definition of equivalence that identifies more games as equivalent is *best response* equivalence. Some notation:

- $N - i$ denotes the set of players other than i;
- C_{-i} denotes the set of all possible combinations of strategies for the players other than i: $C_{-i} = \Pi_{j\in N-i} C_j$;
- Given any $e_{-i} = (e_j)_{j\in N-i}$ in C_{-i} and any d_i in C_i, $(e_{-i}, d_i) \in C$ is a strategy profile;
- For any set Z and any function $f : Z \to \mathbb{R}$, we define $argmax_{y\in Z} f(y) = \{y \in Z \mid f(y) = max_{z\in Z} f(z)\}$;
- Suppose player i believed that some probability distribution $\eta \in \Delta(C_{-i})$ predicted the behavior of the other players in the game, then in game Γ, player i's set of *best responses* to η would be $argmax_{d_i\in C_i} \quad \Sigma_{e_{-i}\in C_{-i}} \quad \eta(e_{-i})u_i(e_{-i}, d_i)$.

Definition 2.13. *Two games Γ and Γ' as defined in the preceding definition are* best-response equivalent *iff the best-response sets always coincide, i.e.* $\forall_{i \in N} \forall_{\eta \in \Delta(C_{-i})}$

$$argmax_{d_i \in C_i} \quad \Sigma_{e_{-i} \in C_{-i}} \eta(e_{-i}) u_i(e_{-i}, d_i) \quad = \quad argmax_{d_i \in C_i} \quad \Sigma_{e_{-i} \in C_{-i}} \eta(e_{-i}) u'_i(e_{-i}, d_i)$$

Example 2.14. *Consider the following two games.*

C_1	C_2 :	x_2	y_2
x_1		*9,9*	*0,8*
y_1		*8,0*	*7,7*

C_1	C_2 :	x_2	y_2
x_1		*1,1*	*0,0*
y_1		*0,0*	*7,7*

The following two games are best-response equivalent because, in each game, each player i's best response would be to choose y_i if he thought that the other player would choose the y-strategy with probability $1/8$ or more. Otherwise, i's best response would be to choose x_i.

This can be verified as follows:
Consider player 1, having two strategies x_1 and y_1. Suppose that $\eta(x_2) = \eta(y_2) = 1/2$. Then

$$\Sigma_{e_2 \in C_2} \eta(e_2) u_1(e_2, x_1) = 1/2 \times 9 + 1/2 \times 0 = 4.5,$$

while

$$\Sigma_{e_2 \in C_2} \eta(e_2) u_1(e_2, y_1) = 1/2 \times 8 + 1/2 \times 7 = 7.5.$$

So

$$argmax_{d_1 \in C_1} \quad \Sigma_{e_2 \in C_2} \eta(e_2) u_1(e_2, d_i) = \{y_1\}.$$

Likewise we have

$$\Sigma_{e_2 \in C_2} \eta(e_2) u'_1(e_2, x_1) = 1/2 \times 1 + 1/2 \times 0 = 0.5,$$
$$\Sigma_{e_2 \in C_2} \eta(e_2) u'_1(e_2, y_1) = 1/2 \times 0 + 1/2 \times 7 = 4.5.$$

This also leads to the preference $\{y_1\}$ for player 1. We can verify the same for the other player and generalize it to an arbitrary probability distribution.

On the other hand, the two games are not fully equivalent, because e.g. each player prefers (x_1, x_2) over (y_1, y_2) in the first game but not in the second game.

2.4 Reduced normal form representations

The idea is to remove superfluous information from games in normal form.

Example 2.15. *Consider the following game.*

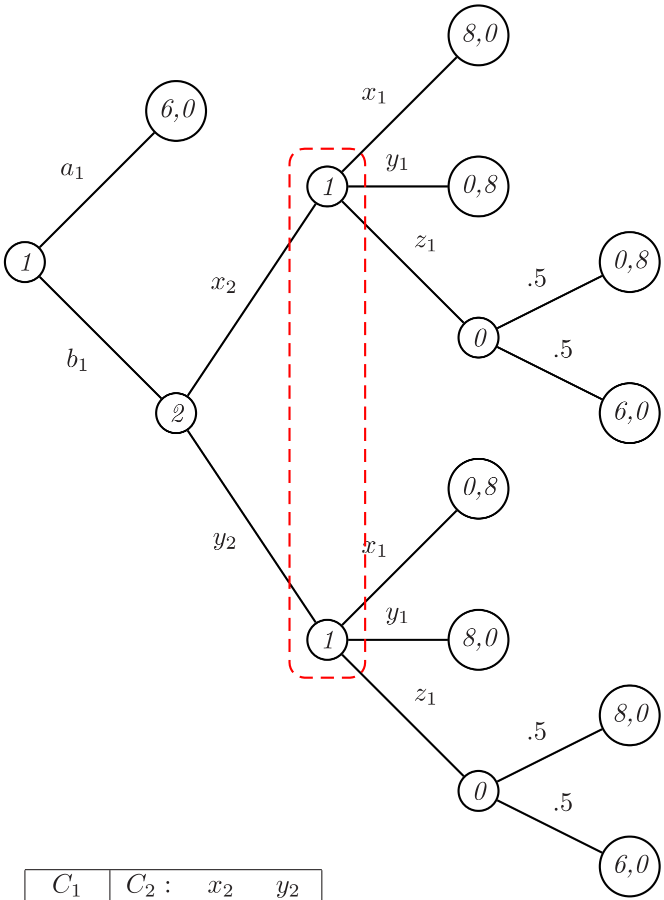

C_1	C_2 :	x_2	y_2
a_1x_1		6,0	6,0
a_1y_1		6,0	6,0
a_1z_1		6,0	6,0
b_1x_1		8,0	0,8
b_1y_1		0,8	8,0
b_1z_1		3,4	7,0

The first three lines express the same information. That is, the utility is independent of the choice between x_1, y_1, and z_1.

Definition 2.16. *Let $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$, let $i \in N$ and $d_i, e_i \in C_i$. We say that d_i and e_i are* payoff equivalent *iff* $\forall_{c_{-i}\in C_{-i}} \forall_{j\in N} \quad u_j(c_{-i}, d_i) = u_j(c_{-i}, e_i)$

Reducing a game in strategic form according to payoff equivalence gives the *purely reduced normal representation.*

Example 2.17. *Below is the purely reduced normal representation of the previous example.*

C_1	C_2 :	x_2	y_2
a_1-		6,0	6,0
b_1x_1		8,0	0,8
b_1y_1		0,8	8,0
b_1z_1		3,4	7,0

Furthermore, we observe that strategy b_1z_1 is a linear combination of strategies a_1- and b_1y_1: $1/2 \times (6,0) + 1/2 \times (0,8) = (3,4)$*, and* $1/2 \times (6,0) + 1/2 \times (8,0) = (7,0)$.

Some notation:

- Strategies in the set C_i are called *pure strategies* for player i;

- Combinations of strategies in the set $\Delta(C_i)$ are called *randomized strategies* for i, i.e. a randomized strategy for i is a probability distribution over its pure strategies.

Definition 2.18. *A strategy d_i in C_i is* randomly redundant *iff there is a probability distribution σ_i in $\Delta(C_i)$ such that $\sigma_i(d_i) = 0$ and*

$$\forall_{c_{-i} \in C_{-i}} \forall_{j \in N} \quad u_j(c_{-i}, d_i) = \Sigma_{e_i \in C_i}\ \sigma_i(e_i) u_j(c_{-i}, e_i).$$

The *fully reduced normal representation* is derived from the purely reduced normal representation by eliminating all strategies that are randomly redundant.

Example 2.19. *Below is the fully reduced normal representation of the previous example.*

C_1	$C_2:$	x_2	y_2
a_1-		*6,0*	*6,0*
b_1x_1		*8,0*	*0,8*
b_1y_1		*0,8*	*8,0*

Notation

$\Delta(Z)$	Set of Probability Distributions
$X \ni x, y, z$	Set of prizes
$\Omega \ni t, r$	Set of states
$L \ni f, g, h$	Set of all lotteries
$f(x\|t) \in \Delta(X)$ (fixed t)	lottery f
$\Xi = \{S \mid S \subset \Omega, S \neq \emptyset\}$	Set of world events
$f \gtrsim_S g$, $f \sim_S g$, $f >_S g$	Preferences in lotteries
$\alpha f + (1-\alpha)g$	lottery $\alpha f(x\|t) + (1-\alpha)g(x\|t)$
$[x]$	lottery that always gives prize x for sure.
$p(t\|S) \in \Delta(\Omega)$ (fixed S)	Conditional-probability function
$u : X \times \Omega \to \mathbb{R}$	Utility function
$E_p(u(f)\|S) = \sum_{t \in S} p(t\|S) \sum_{x \in X} u(x,t) f(x\|t)$	Expected Utility Value