Chapter 11, sections 1, 4, and 5

Extensive form games with imperfect information

Generalize extensive form games with perfect information in two directions:

(1) there may be moves by chance/nature

(2) players may be imperfectly informed about earlier choices

Histories that a player cannot distinguish are gathered in an information set.

Def. 200.1 An *extensive form game*

$$\Gamma = \langle N, H, P, \underbrace{f_c, (\mathcal{I}_i)_{i \in N},}_{\text{new stuff}} (\succsim_i)_{i \in N} \rangle$$

consists of

☒ nonempty, finite set N of *players*,

☒ set H of *histories*, summarizing the sequence of actions/chance moves so far.

- same restrictions as before in def. 89.1,
- $Z \subseteq H$ set of *terminal histories,*
- actions after nonterminal history $h \in H$ are $A(h) = \{a \mid (h, a) \in H\}$,

⊠ *player function* $P : H \setminus Z \to N \cup \{c\}$ assigning to each nonterminal history a player/chance whose turn it is to take an action,

⊠ f_c assigns to each history $h \in H$ with $P(h) = c$ a prob measure $f_c(\cdot \mid h)$ over actions $A(h)$, indep of every other such measure $f_c(\cdot \mid h')$. Roughly (finite case), $f_c(a \mid h)$ is the prob that chance chooses a.

⊠ for each player $i \in N$, $\mathcal{I}_i$ is a partition of $\{h \in H \mid P(h) = i\}$ such that

$$\forall I_i \in \mathcal{I}_i, \forall h, h' \in I_i : \qquad A(h) = A(h'). \tag{1}$$

(Same set of actions in each history of an information set: otherwise i could distinguish the histories!)

- $\mathcal{I}_i$ is the *information partition*, $I_i \in \mathcal{I}_i$ is an *information set*.

- idea: i cannot distinguish histories h, h' in the same information set.

- define $P(I_i) = i$ and $A(I_i) = A(h)$ for arbitrary $h \in I_i$. Well-defined by (1)!

☒ for each player $i \in N$ a *preference relation* $\succsim_i$ over terminal histories. Often assumed to satisfy the vNM axioms, in which case we represent them by a Bernouilli function u_i on Z.

Remark: extensive form games with perfect information

$$\Gamma = \langle N, H, P, (\succsim_i)_{i \in N} \rangle$$

are embedded as a special case: no chance moves and each information set is a singleton (only one history).

Notational convention: when drawing game trees, information sets are indicated by ovals/dotted lines.

Fig. 212.1 An extensive form game with chance move

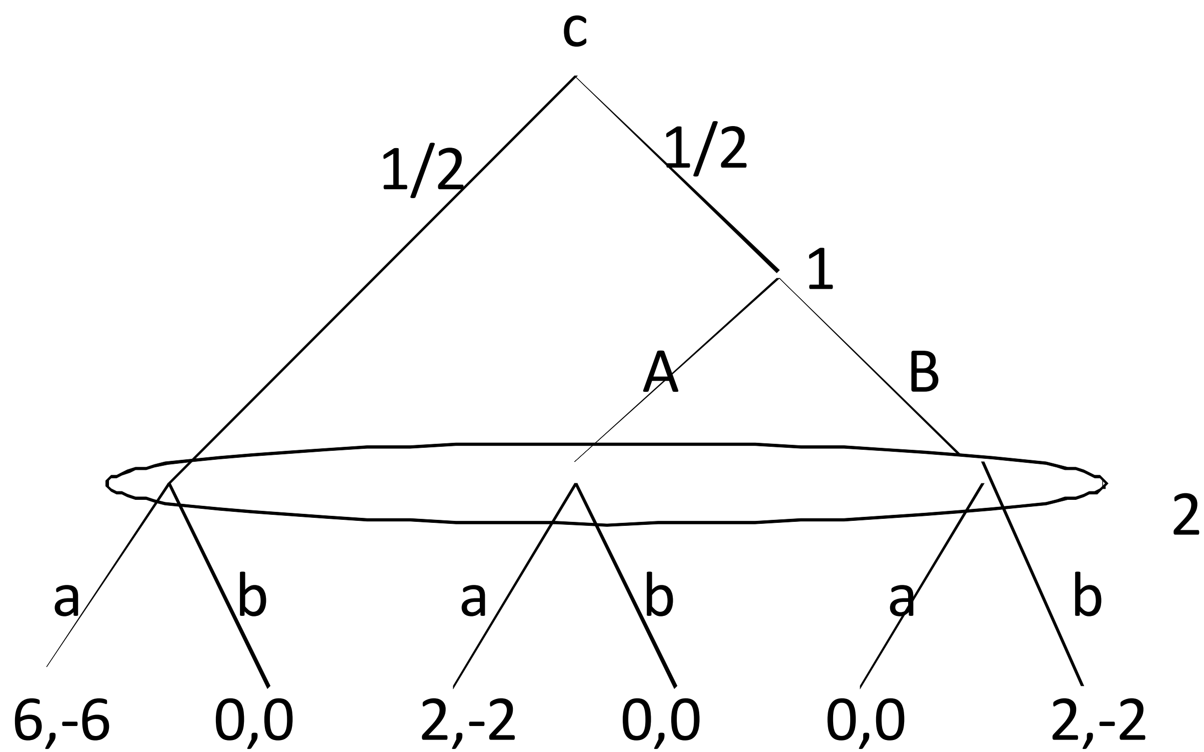

Strategic form:

	a	b
A	$4, -4$	$0, 0$
B	$3, -3$	$1, -1$

(Im)perfect recall

Do you remember what you knew in the past?

Example 1. Absentminded driver

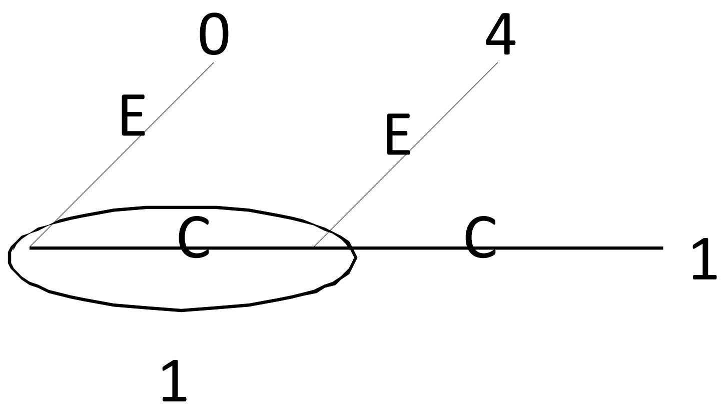

Two crossings on your way home. You need to (C)ontinue on the first, (E)xit on the second. But you can't recall *whether* you already passed a crossing.

Example 2. The player forgot what he decided in the initial node.

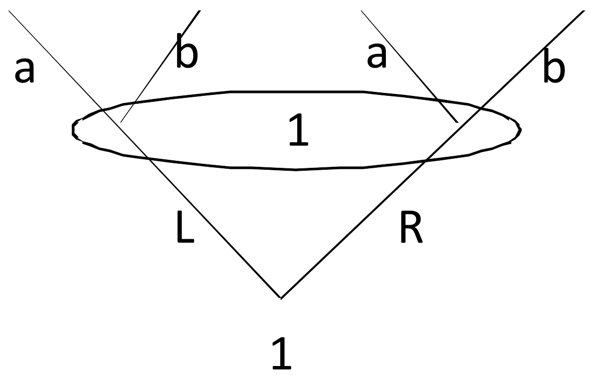

Example 3. Pl 1 knew the chance move, but forgot it.

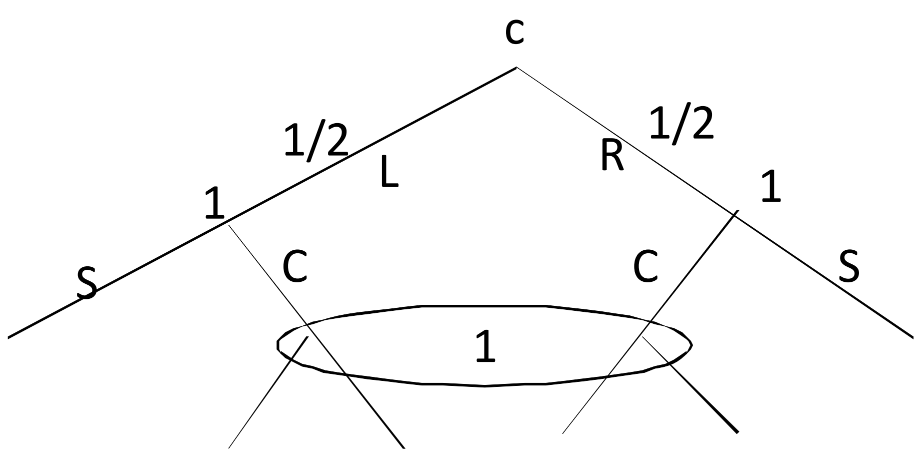

Experience and perfect recall

Consider a history $h \in H$ with $P(h) = i$. The *experience* $X_i(h)$ along history h lists, in chronological order:

(1) which information sets I_i player i encountered

(2) what action i chose there

Def. 203.3 Extensive form game $\Gamma = \langle N, H, P, f_c, (\mathcal{I}_i)_{i \in N}, (\succsim_i)_{i \in N} \rangle$ has *perfect recall* if

$$\forall i \in N, \forall I_i \in \mathcal{I}_i, \forall h, h' \in I_i : \qquad X_i(h) = X_i(h').$$

You have the same experience in each history of your information set: you recall exactly what sequence of events got you there!

Example 1 (cont)

☒ Only one information set, $\{\emptyset, C\}$,

☒ $X_1(\emptyset) = (\{\emptyset, C\})$, but

☒ $X_1(C) = (\underbrace{\{\emptyset, C\}}_{\text{1's first info set}}, \underbrace{C}_{\text{choice there}}, \underbrace{\{\emptyset, C\}}_{\text{resulting info set}})$

☒ $X_1(\emptyset) \neq X_1(C)$: imperfect recall!

Example 2 (cont)

☒ Consider information set $\{L, R\}$.

☒ $X_1(L) = (\{L, R\}, L, \{L, R\})$

☒ $X_1(R) = (\{L, R\}, R, \{L, R\})$

☒ $X_1(L) \neq X_1(R)$: imperfect recall!

Example 3 (cont)

☒ Consider information set $\{(L,C),(R,C)\}$.

☒ $X_1((L,C)) = (\underbrace{\{L\}}_{\text{1's first info set}}, \underbrace{C}_{\text{choice there}}, \underbrace{\{(L,C),(R,C)\}}_{\text{resulting info set}})$

☒ $X_1((R,C)) = (\{R\}, C, \{(L,C),(R,C)\})$

☒ $X_1((L,C)) \neq X_1((R,C))$: imperfect recall!

Exc. 203.2 Can you define experience formally?

☒ If $h = \emptyset$, $X_i(\emptyset)$ consists of the information set containing the initial node.

☒ Consider a history $h = (a^1, \ldots, a^k) \in H$ with $P(h) = i$.

☒ Write $h^0 = \emptyset$ and for $r = 1, \ldots, k : h^r = (a^1, \ldots, a^r)$.

☒ Steps after which i acts: $R(i) = \{r \in \{0, \ldots, k\} \mid P(h^r) = i\}$.

☒ Corresponding info sets: for each $r \in R(i)$, let $I_i^r \in \mathcal{I}_i$ be the info set containing h^r.

☒ Order the elements of $R(i)$ in increasing order: $r_1 < \cdots < r_\ell$.

☒ $X_i(h) = \{ \underbrace{I_i^{r_1}}_{1^{st} \text{ info set of } i}, \underbrace{a^{r_1+1}}_{\text{what } i \text{ does there}}, \underbrace{I_i^{r_2}, a^{r_2+1}, \ldots, I_i^{r_\ell - 1}, a^{r_\ell}, I_i^{r_\ell}}_{\text{etc.}} \}.$

Three types of strategies

☒ Consider an extensive form game $\Gamma = \langle N, H, P, f_c, (\mathcal{I}_i)_{i \in N}, (\succsim_i)_{i \in N} \rangle$.

☒ **Def. 203.1:** a *pure strategy* of pl $i \in N$ is a function s_i that assigns to each information set $I_i \in \mathcal{I}_i$ of pl i a feasible action $s_i(I_i) \in A(I_i)$.

☒ **Def. 212.1**: a *mixed strategy* of pl $i \in N$ is a probability distribution σ_i over the pure strategies of pl i.

"Global randomization" at the beginning of the game.

☒ **Def. 212.1:** a *behavioral strategy* of pl. $i \in N$ is a collection $\beta_i = (\beta_i(I_i))_{I_i \in \mathcal{I}_i}$ of (indep) probability measures over the actions in i's information sets: $\beta_i(I_i)$ is a prob meas over $A(I_i)$.

"Local randomization" as play proceeds.

Other terminology/notational conventions:

☒ let $h \in I_i \in \mathcal{I}_i$ and $a \in A(h)$. The prob assigned to a is denoted by $\beta_i(h)(a)$ or $\beta(I_i)(a)$.

Outcomes

Each profile $\sigma = (\sigma_i)_{i \in N}$ of mixed strategies induces an *outcome* $O(\sigma)$, a prob measure over terminal outcomes.

How to compute $O(\sigma)$ in finite games?

⊠ Let $h = (a^1, \ldots, a^k)$ be a history.

⊠ Pure strategy s_i of player i is *consistent with* h if i chooses the actions described by h: for each initial segment $(a^1, \ldots, a^\ell)$ with $P(a^1, \ldots, a^\ell) = i$:

$$s_i(a^1, \ldots, a^\ell) = a^{\ell+1}.$$

⊠ The prob of choosing a pure strategy s_i that is consistent with h is

$$\pi_i(h) = \sum_{s_i \text{ consistent with } h} \sigma_i(s_i).$$

(similar for nature, whose behavior is given by f_c).

☒ Prob of reaching terminal node h is

$$\prod_{i \in N \cup \{c\}} \pi_i(h).$$

Similarly, each profile $\beta = (\beta_i)_{i \in N}$ of behavioral strategies induces an *outcome* $O(\beta)$, a prob measure over terminal outcomes.

How to compute $O(\beta)$ in finite games?

⊠ Let $h = (a^1, \ldots, a^k)$ be a terminal history.

⊠ The prob of reaching terminal node $h = (a^1, \ldots, a^k)$ is

$$\prod_{\ell=0}^{k-1} \beta_{P(a^1,\ldots,a^\ell)}(a^1, \ldots, a^\ell)(a^{\ell+1}).$$

Does the distinction between mixed/behavioral strategies matter?

Not if they generate the same probability measures over terminal histories.

Strategies σ_i and β_i of pl $i \in N$ are *outcome-equivalent* if — given the pure strategies of the remaining players — they give rise to the same outcome (as in lecture 9):

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

Prop. 214.1 (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. strategy β_i, assign to pure strategy s_i the probability

$$\sigma_i(s_i) = \prod_{I_i \in \mathcal{I}_i} \beta_i(I_i)(s_i(I_i)).$$

[Intuition: s_i selects action $s_i(I_i)$ in info set I_i. How likely is that?]

(b) Given mixed strategy σ_i, consider an info set $I_i \in \mathcal{I}_i$ of pl i and an action $a \in A(I_i)$. How should we define $\beta_i(I_i)(a)$?

Consider history $h \in I_i$. Probability of choosing consistent with h is $\pi_i(h)$.

[Perfect recall: $\pi_i(h) = \pi_i(h')$ for all $h, h' \in I_i$].

Define

$$\beta_i(I_i)(a) = \frac{\pi_i(h, a)}{\pi_i(h)} \text{ if } \pi_i(h) > 0 \text{ (and arbitrarily otherwise).}$$

[Intuition: conditional on being consistent with h, how likely is i to choose a?]

Example 1 (cont)

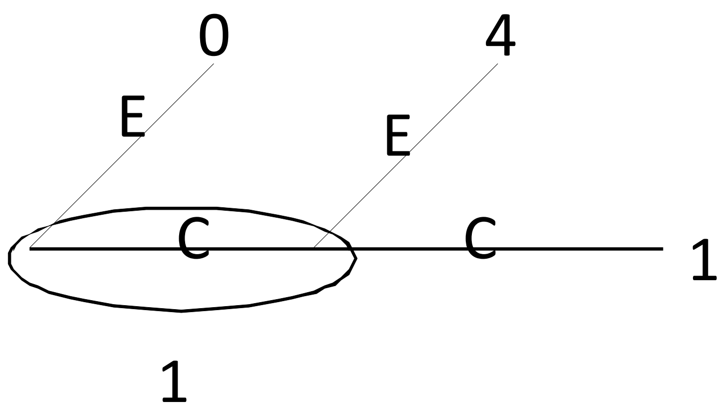

Pure strategies?

Mixed strategies?

Behavioral strategies?

Outcome equivalence?

☒ Pure: Only one information set $I_1 = \{\emptyset, C\}$, two actions: C and E. So two pure strategies, with associated payoffs

$$\begin{array}{ll} C & 1 \\ E & 0 \end{array}$$

☒ Mixed: let $p = \sigma_1(C)$ be the prob of choosing pure strategy C (and $1 - p$ the prob of choosing E). Expected payoff: p

☒ Behavioral: let $q = \beta_1(I_1)(C)$ 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?)

Exc. 216.1

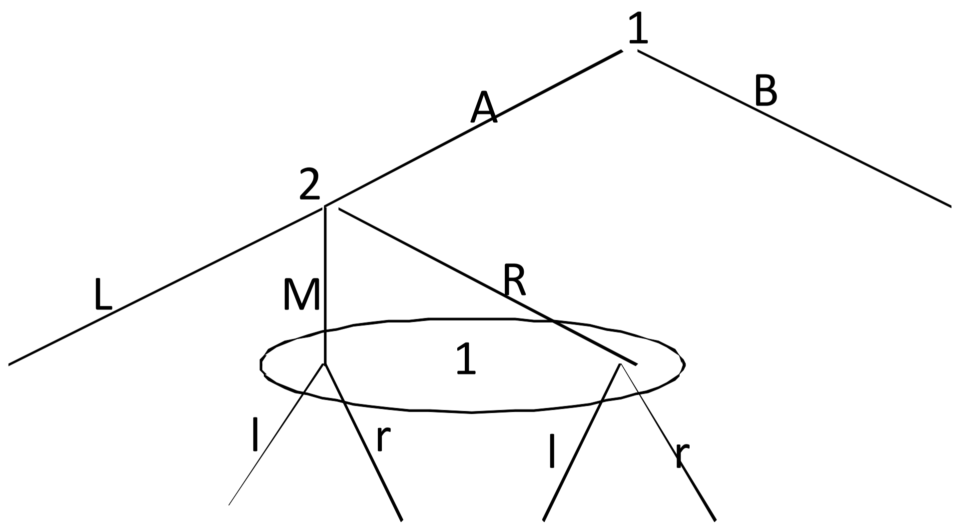

Which behavioral strategy is outcome-equivalent with mixed strategy

$$\sigma_1(A\ell) = 5/10, \sigma_1(Ar) = 0, \sigma_1(B\ell) = 1/10, \sigma_1(Br) = 4/10?$$

☒ $\beta_1(\emptyset)(A) = \pi_1(A)/\pi_1(\emptyset) = (5/10)/1 = 1/2$ and $\beta_1(\emptyset)(B) = 1/2$.

☒ $\beta_1(\{(A, M), (A, R)\})(\ell) = \pi_1((A, M, \ell))/\pi_1(A, M) = (5/10)/(5/10) = 1$ and $\beta_1(\{(A, M), (A, R)\})(r) = 0$.

Nash equilibrium

☒ As in the case of extensive form games with perfect information, we can construct a corresponding strategic game.

☒ A *Nash equilibrium in mixed strategies* is a mixed-strategy profile $\sigma^* = (\sigma_i^*)_{i \in N}$ such that no player can profitably deviate to another mixed strategy:

for all $i \in N$ and all mixed strategies σ_i of pl i : $\qquad O(\sigma^*) \succsim_i O(\sigma^*_{-i}, \sigma_i)$.

This is just a mixed Nash equilibrium of the corresponding strategic game.

☒ *Nash equilibrium in behavioral strategies* is defined likewise.

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

Proof: Finite extensive form game results in finite strategic game, which has a Nash eq. in mixed strategies (lecture 1). By outcome-equivalence, we can construct a Nash equilibrium in behavioral strategies.