Sequential equilibrium (Ch. 12, sec 1, 2 (until p. 227), 5)

Recall: subgame perfect eq 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.

An obvious translation to extensive games with imperfect information would be to require best responses in each information set.

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

Fig. 220.1

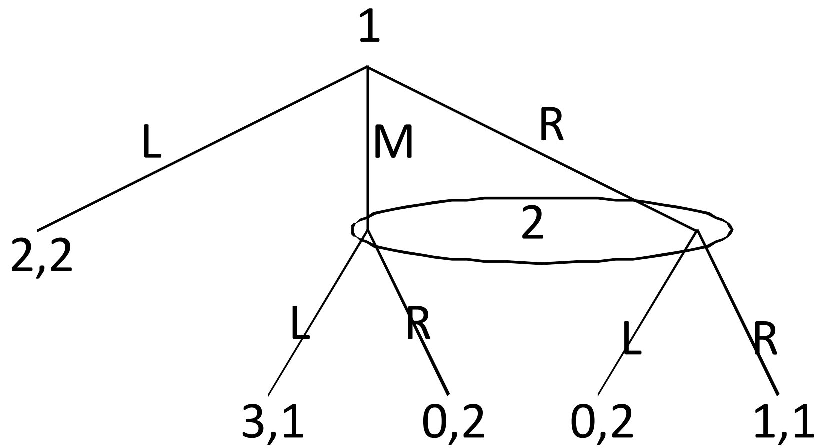

Beliefs affect strategies: Pl. 2 in info set $\{M, R\}$: L is a best response if and only if pl. 2 assigns prob at most 1/2 to being in history M.

Strategies affect 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 history R than in history M. Bayes' Law requires that he assigns prob. 1/3 to M and 2/3 to R.

Q: "sensible" beliefs if pl. 1 chooses L with prob 1?

No obvious answer. We'll see different approaches:

- today: consistency
- next lecture: weak consistency/structural consistency

The equilibrium notion we introduce today involves an interplay between beliefs and strategies.

We restrict attention to extensive form games

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

with perfect recall where each information set has *finitely many* histories.

Def. 222.1: An *assessment* (β, μ) consists of

☒ a profile $\beta = (\beta_i)_{i \in N}$ of behavioral strategies,

☒ a *belief system* μ that assigns to each information set a probability distribution over its histories.

If I is an information set of pl. i and $h \in I$ a history in that info set, then $\mu(I)(h)$ denotes the probability that pl. i assigns to history h if he is in info set I.

Outcomes

This helps us to define rational choices in an information set. Given assessment (β, μ) and information set I, we define the *outcome* $O(\beta, \mu \mid I)$ — the prob distr over terminal nodes — of (β, μ) conditional on reaching I as follows:

☒ Let $h^* = (a^1, \ldots, a^K)$ be a terminal history.

☒ If I does not contain a subhistory $h = (a^1, \ldots, a^L)$ $(L < K)$ of h^* (i.e., information set I rules out history h^*), then $O(\beta, \mu \mid I)(h^*) = 0$.

☒ If I does contain a subhistory $h = (a^1, \ldots, a^L)$ $(L < K)$ of h^*, then

$$O(\beta, \mu \mid I)(h^*) = \underbrace{\mu(I)(h)}_{\text{prob of } h \text{ given } I} \cdot \underbrace{\prod_{\ell=L}^{K-1} \beta_{P(a^1,\ldots,a^\ell)}(a^{\ell+1})}_{\text{prob that players proceed with } (a^{L+1},\ldots,a^K)}$$

Sequential rationality

Def. 224.1: An assessment (β, μ) is *sequentially rational* if each player in each information set chooses a best response given his beliefs and the strategies of the other players:

$$\forall i \in N, \forall I_i \in \mathcal{I}_i, \forall \beta_i' : \qquad O(\beta, \mu \mid I_i) \succsim_i O((\beta_{-i}, \beta_i'), \mu \mid I_i).$$

Notice: sequential rationality of (β, μ) does not imply that β is a Nash equilibrium: there are no restrictions on beliefs, so players might be best-replying against their beliefs, but not against the strategies that are actually played!

In Fig. 220.1, the pure strategy profile (M, L) is sequentially rational if the belief system (erroneously) has player 2 assigning prob 1 to being in history R whenever his info set is reached.

Belief restrictions: consistency

☒ If all info sets are reached with positive probability, we know how to formulate reasonable beliefs: use Bayes' Law.

☒ A profile β of behavioral strategies is *completely mixed* if it assigns positive probability to each action in each information set.

☒ What if some info sets are reached with zero probability? One potential resolution is to consider this as a "limit case" of completely mixed strategy profiles:

Def. 224.2: An assessment (β, μ) is *consistent* if there is a sequence $(\beta^k, \mu^k)_{k=1}^{\infty}$ of assessments such that:

☒ for each $k \in \mathbb{N}$: β^k is completely mixed,

☒ for each $k \in \mathbb{N}$: μ^k follows from β^k using Bayes' Law,

☒ $(\beta^k, \mu^k) \to (\beta, \mu)$.

Def. 225.1: An assessment (β, μ) is a *sequential equilibrium* if it is

☒ sequentially rational and

☒ consistent.

Prop. In a finite extensive game with perfect information (singleton information sets), a behavioral strategy profile is a subgame perfect equilibrium if and only if it is a sequential equilibrium (with the only feasible belief system that assigns prob 1 to the single history in each info set).

Prop. If (β, μ) is a sequential equilibrium, then β is a Nash equilibrium.

Proof sketch: Suppose β were not a NE. Let i have profitable deviation β_i'. Then somewhere on the equilibrium path (i.e., in a history reached with positive probability) there is an info set where i has a profitable deviation, contradicting sequential rationality (and the correctness of beliefs).

Example. Determine all sequential equilibria of the game below.

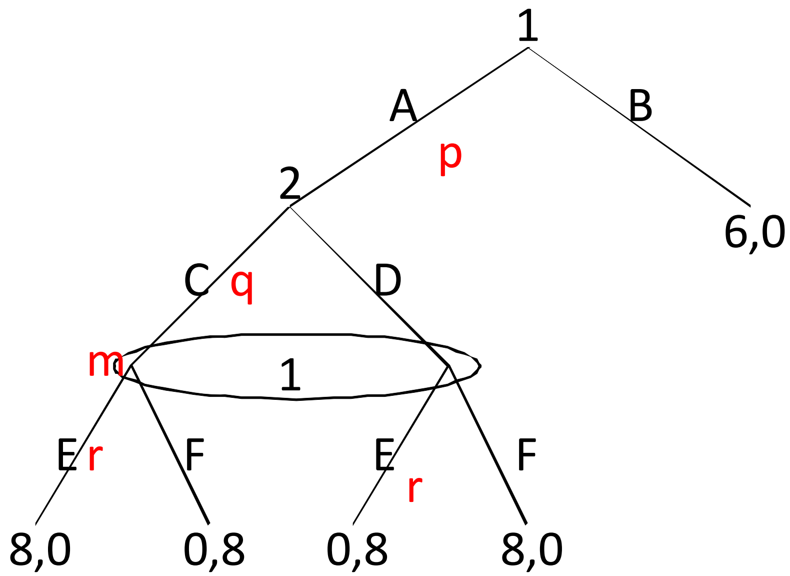

Intuition: what should it be? pl. 1 chooses between sure payoff $(6, 0)$ or the bimatrix game

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

Notation: for notational convenience, we characterize a behavioral strategy $\beta = (\beta_1, \beta_2)$ by

$$p = \beta_1(\emptyset)(A), q = \beta_2(A)(C), r = \beta_1(\{(A, C), (A, D)\})(E)$$

and the belief system μ by

$$m = \underbrace{\mu(\{(A, C), (A, D)\})(A, C)}_{\text{prob assigned to } (A,C) \text{ in 2's info set}}$$

Consistency: A completely mixed beh str profile has $p, q, r \in (0, 1)$. Bayes' rule then implies that

$$m = \frac{pq}{pq + p(1-q)} = q.$$

So for each consistent (β, μ), it follows that $m = q$.

Sequential rationality: which (β, μ) with $m = q$ are sequentially rational?

Distinguish 3 cases:

☒ If $q = 0$, then $m = 0$, so $r = 0$ is seq. rat. But if $r = 0$, then $q = 0$ is not seq. rat. Contradiction!

☒ If $q = 1$, then $m = 1$, so $r = 1$ is seq. rat. But if $r = 1$, then $q = 1$ is not seq. rat. Contradiction!

☒If $q \in (0, 1)$, seq. rat. in info set $\{A\}$ means 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, the expected payoff is

$$m[8r] + (1 - m)[8(1 - r)] \underbrace{=}_{m=q} 8qr + 8(1 - q)(1 - r) = 8 - 8q + 8r(2q - 1).$$

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: the unique sequential equilibrium has

$$p = 0, q = m = r = 1/2.$$

Existence of sequential equilibria

Requires a bit of a detour using a new equilibrium concept for strategic form games and thereafter for extensive form games:

Trembling-hand perfect equilibria in strategic games

Requires a best response not only to the other players' strategies, but also to nearby ones

Def. 248.1: A *trembling-hand perfect equilibrium* of a finite strategic game $\langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a mixed strategy profile $\sigma = (\sigma_i)_{i \in N}$ with the property that there is a sequence $(\sigma^k)_{k=1}^{\infty}$ of completely mixed strategy profiles (pos prob on each pure strategy!) such that

(1) $\sigma^k \to \sigma$

(2) for each pl. $i \in N$ and each $k \in \mathbb{N}$: σ_i is a best response to σ^k_{-i}.

Prop. 248.2 and 249.1: Consider a finite strategic game

$$G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle.$$

(a) G has at least one trembling-hand perfect equilibrium.

(b) Each trembling-hand perfect equilibrium is a Nash equilibrium.

(c) If $\sigma = (\sigma_i)_{i \in N}$ is a trembling-hand perfect equilibrium, then no σ_i is weakly dominated.

(d) If G has exactly two players and $\sigma = (\sigma_1, \sigma_2)$ is a Nash equilibrium where no σ_i is weakly dominated, then σ is trembling-hand perfect.

Proof sketch:

(a) For $\varepsilon > 0$ sufficiently small, consider the game $G(\varepsilon)$ where the strategy space of each player i is restricted to

$$\{\sigma_i : A_i \to [0,1] \mid \sigma_i(a_i) \geq \varepsilon \text{ for each } a_i \in A_i \text{ and } \sum_{a_i} \sigma_i(a_i) = 1\}.$$

(Each action chosen with prob at least ε). By lecture 1, it has a Nash equilibrium σ^ε. Now let $(\varepsilon^k)_{k=1}^\infty$ converge to zero and let σ^k be a Nash eq of $G(\varepsilon^k)$. Since this gives us a seq in the compact strategy space of the original game G, we can consider a convergent subsequence — w.l.o.g. the sequence $(\sigma^k)_{k=1}^\infty$ itself — with limit σ: requirement (1) holds. Requirement (2) holds by continuity of payoffs and σ^k being a Nash eq of $G(\varepsilon^k)$.

(b) If σ_i is a best response to a sequence of nearby σ_{-i}^k, it is a best response to their limit by continuity of the payoff function.

(c) Consider a sequence $(\sigma^k)_{k=1}^{\infty}$ that makes σ trembling hand perfect. Suppose σ_i is weakly dominated by τ_i :

$$\begin{aligned} \forall a_{-i} \in A_{-i} : \quad & u_i(\tau_i, a_{-i}) \geq u_i(\sigma_i, a_{-i}), \\ \exists a_{-i} \in A_{-i} : \quad & u_i(\tau_i, a_{-i}) > u_i(\sigma_i, a_{-i}), \end{aligned}$$

Then for each $k \in \mathbb{N}$:

$$\underbrace{u_i(\tau_i, \sigma_{-i}^k)}_{\substack{\text{by def a convex combination of } u_i(\tau_i, a_{-i}), \\ \text{with all weights positive, since } \sigma^k \text{ is completely} \\ \text{mixed}}} > u_i(\sigma_i, \sigma_{-i}^k),$$

contradicting requirement that σ_i is a best response to σ_{-i}^k.

(d) For each pl. i, using a separating hyperplane argument, it can be shown that σ_i is a best reply to some completely mixed strategy τ_{-i} of the other player. The sequence

$$\left(\frac{k-1}{k}\right)\sigma + \frac{1}{k}\tau$$

converges to σ and σ_1 is a best response to $\left(\frac{k-1}{k}\right)\sigma_2 + \frac{1}{k}\tau_2$ since

$$u_1(\sigma_1, \left(\frac{k-1}{k}\right)\sigma_2 + \frac{1}{k}\tau_2) = \left(\frac{k-1}{k}\right)u_1(\sigma) + \frac{1}{k}u_1(\sigma_1, \tau_2)$$

cannot be improved upon: σ is a Nash equilibrium (no improvements in first summand) and σ_1 is a best response to τ_2 (no improvements in second summand).

Trembling hand perfect equilibria in extensive form games

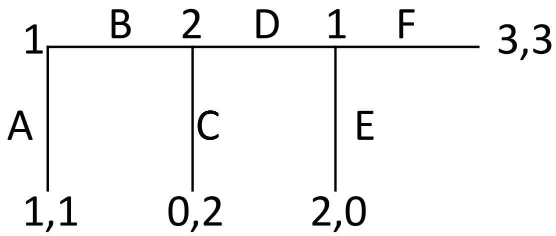

Strategic form:

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

Unique subgame perfect equilibrium: (BF, D)

In the strategic form, this is a trembling-hand perfect equilibrium, but so is — for instance — (AE, C).

Why? Firstly, it is an undominated Nash equilibrium in a two-player game! Secondly (closer to the original definition):

AE is a best response to any mixture over L and R with large probability on C.

C is a best response to any mixture over AE, AF, BE, and BF as long as the prob of AE is close to one and the prob of BE is large w.r.t. BF.

Roughly: (AE, C) is robust to trembles if they are defined over pure strategies, but not if trembles take place in each history/information set separately: then pl. 1 would rather choose F than E.

Consider the *agent-strategic form*:

☒ there is one player for each information set in the game: each player of the extensive form game is split up into agents, one agent for each of information set.

☒ all agents of a player i have the same preferences as i.

Formally, the player set is $\{(i, I_i) \mid i \in N, I_i \in \mathcal{I}_i\}$ and player (i, I_i) has the same preferences over terminal histories as pl. i.

The game above has 2 agents for player 1 (who has 2 info sets, namely $\{\emptyset\}$ and $\{(B, D)\}$) and only one agent for player 2. The corresponding agent strategic form is:

	C	D
A	$1,1,1$	$1,1,1$
B	$0,0,2$	$2,2,0$
	E	

	C	D
A	$1,1,1$	$1,1,1$
B	$0,0,2$	$3,3,3$
	F	

Notice: a mixed-strategy profile in the agent strategic form corresponds with a behavioral strategy profile in the extensive form, where $\beta_i(I_i)$ is the mixed strategy of player i's agent in information set I_i.

Def. 251.1: A *trembling-hand perfect equilibrium of a finite extensive form game* is a behavioral strategy profile that corresponds to a trembling hand perfect equilibrium of the agent strategic form.

Prop. 251.2 and 253.2: Consider a finite extensive form game with perfect recall Γ.

(a) For each trembling hand perfect equilibrium β of Γ there is a belief system μ such that (β, μ) is a sequential equilibrium.

(b) Γ has at least one sequential equilibrium.

Proof sketch:

(a) Let $(\beta^k)_{k=1}^{\infty}$ be a sequence that makes β trembling hand perfect.

Since β^k is completely mixed, we can define a corresponding belief system μ^k using Bayes' rule.

Let $\mu^k \to \mu$ (taking subsequences if necessary).

It can be shown that (β, μ) is a sequential equilibrium.

(b) A finite extensive form game yields a finite agent strategic form game.

The latter has a trembling-hand perfect equilibrium, which can be extended to a sequential equilibrium according to (a).