Introduction to Game Theory: Sequential Moves
Lecture 10, *Experimental Econ. & Psychology*

Pierpaolo Battigalli
Bocconi University

8 October 2020

Abstract

In sequential game forms, one or more moves are played in sequence, but there may be also simultaneous moves at some stage. In sequential game forms with **perfect information** moves are never simultaneous and players perfectly observe past moves. Such games are the easiest to represent with game trees. Sequences of (updated) beliefs may be as important as sequences of actions. Hence, we are interested in how beliefs change as information accrues to players. From the perspective of psychological game theory, also a static game may have interesting dynamics of beliefs, because the *terminal beliefs* players hold after the play may matter for psychological reasons, so that it may be necessary to try to anticipate how the terminal beliefs of others are affected by one's own actions. We focus for simplicity on games with (at most) two stages, allowing for the possibility of simultaneous moves in either stage.

- We expand our analysis of games to those that may have some sequential moves. How information accrues to players is key.

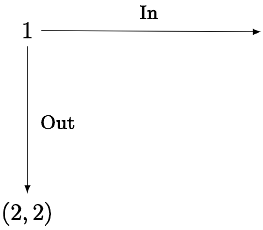

$1\backslash 2$	B_2	S_2
B_1	$3, 1$	$0, 0$
S_1	$0, 0$	$1, 3$

- **Example:** *Battle of the Sexes with Outside Option.* In the Battle of the Sexes (BoS, a static game) players would like to coordinate (e.g., both go to Ballet, or both go to Stadium), but have conflicting interests in how to coordinate. In the BoSOO (pictured), pl. 1 can take an **Out**side Option or go **In** and play the BoS. We assume that chosen actions are commonly observed.

- The old formalism "pretends" that there is *only one active player* (possibly, chance) *in each situation*. Simultaneous moves are "simulated" by imposing an arbitrary sequence and assuming that late movers do not observe the choice of early movers.
- **Example:** *Tree representation of BoSOO*. In the BoS "subgame" pl. 1 moves first, but pl. 2 cannot observe 1's choice.

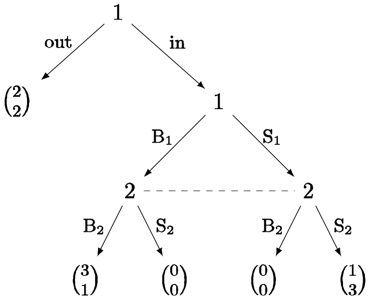

- For traditional GT, such distorted sequential representation of simultaneous moves is innocuous. Experiments show it is not: e.g., in the BoS early movers tend to be advantaged even if their choice is not observed. *Once neglected details may matter for psychological reasons.* Thus, the new formalism strives for a more accurate representation.
- **Example:** A faithful representation of BoSOO does not prevent tree-like pictures:

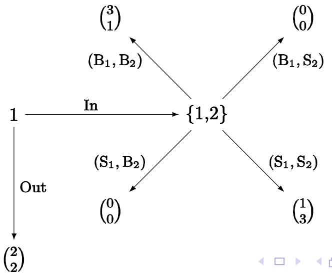

Sequential game forms: the game tree

- $I_0 = I \cup \{0\}$, *finite* player set, including the chance pl. 0.
- $\bar{H}$, *finite* set of possible sequences of action profiles (**histories**) $h = \left(a^k\right)_{k=1}^{\ell}$ including the **empty sequence** $\varnothing$ (**root**) s.t., for every $h \in \bar{H}$, every **prefix** (initial subsequence) of h is in $\bar{H}$: $(\bar{H}, \prec)$—where $\prec$ is the "prefix of" relation—is a *tree* with nodes $h \in \bar{H}$.
  - Z=set of **terminal** histories/nodes (game over); H=set of **non-terminal** histories/nodes (including root $\varnothing$).
  - $\iota : H \rightrightarrows I_0$ is the **active-players correspondence**: $\iota(h)$=set of active players given h. $H_i = \{h : i \in \iota(h)\}$=nodes where i is active.
  - $A(h) = \times_{i \in \iota(h)} A_i(h)$—s.t. $\forall a$, $a \in A(h) \Leftrightarrow (h, a) \in \bar{H}$—is the set of possible action profiles given h.
- **Example:** BoSOO has $\bar{H} = H \cup Z$ with
  - $H = \{\varnothing, (\mathrm{In})\}$,
  - $Z = \{(\mathrm{Out}), (\mathrm{In}, (\mathrm{B}_1, \mathrm{B}_2)), (\mathrm{In}, (\mathrm{B}_1, \mathrm{S}_2)), ...\}$ (5 elements),
  - $\iota(\varnothing) = \{1\}$, $\iota(\mathrm{In}) = \{1, 2\}$,
  - $A(\varnothing) = \{\mathrm{In}, \mathrm{Out}\}$, $A_i(\mathrm{In}) = \{\mathrm{B}_i, \mathrm{S}_i\}$ ($i = 1, 2$).

Sequential game forms: chance, information, payoffs

- **Chance probability function:** $p_0 = (p_0(\cdot|h))_{h\in H_0}$, with $p_0(\cdot|h) \in \Delta(A_0(h))$ specifies the (objective) probabilities of chance moves.
- *For simplicity*, we assume here that *active players perfectly observe earlier choices*, and we do not represent the non-terminal information of inactive players (not essential for what we study).
- The **terminal information** of each $i \in I$ is given by a *partition* $\mathcal{P}_i$ of Z ($\mathcal{P}_i(z)$ denotes the cell containing z).
- For each $i \in I$, $\pi_i : Z \to Y_i$ is the **material** (e.g., monetary, $Y_i \subseteq$€$\mathbb{R}$) **payoff** function of i.
- **Example:** BoSOO has no chance moves and
  - $\mathcal{P}_i(z) = \{z\}$ for each i and all z (perfect terminal information);
  - $\pi_1(\mathrm{Out}) = 2$, $\pi_1(\mathrm{In}, (\mathrm{B}_1, \mathrm{B}_2)) = 3$, $\pi_1(\mathrm{In}, (\mathrm{B}_1, \mathrm{S}_2)) = 0$, etc.
- Traditional GT: adding utility functions $(v_i : \times_{j\in I} Y_j \to \mathbb{R})_{i\in I}$ we obtain a **game**. Let $u_i = v_i \circ \pi : Z \to \mathbb{R}$.

Example: a reporting game form

- (*Implemented in the lab to study deception*) Initial die roll (where *face-6* of the die *counts* 0). One active (real) player, who privately observes the realization x and then reports a number y: she can lie! A passive player with constant payoff observes (only) the report. The payoff of the active player is equal to her report. (See fig. G_{10} in Battigalli & Dufwenberg 2020.) With this:
- $I_0 = I \cup \{0\} = \{0, 1, 2\}$;
- $H = \{\varnothing\} \cup \{0, ..., 5\}$, $Z = \{0, ..., 5\} \times \{0, ..., 5\}$;
- $\iota(\varnothing) = \{0\}$, $\iota(x) = \{1\}$;
- $A_0(\varnothing) = \{0, ..., 5\}$, $\forall x \in A_0(\varnothing)$, $A_1(x) = \{0, ..., 5\}$;
- $\forall x \in A_0(\varnothing)$, $p_0(x|\varnothing) = \frac{1}{6}$;
- $\forall (x, y) \in Z$, $\mathcal{P}_1(x, y) = \{(x, y)\}$, $\mathcal{P}_2(x, y) = \{0, ..., 5\} \times \{y\}$;
- $\forall (x, y) \in Z$, $\pi_1(x, y) = y$, $\pi_2(x, y) = \text{const.}$

First-order beliefs

Given each $h \in H \cup \mathcal{P}_i$, (real) pl. i has (conditional) 1^{st}-order belief $\alpha_i(\cdot|h) \in \Delta(Z)$ about the play. Systems $\alpha_i = (\alpha_i(\cdot|h))_{h \in H \cup \mathcal{P}_i}$ describe i's *initial* belief about paths, how i would update or revise her beliefs, *and include her terminal beliefs* (this makes the analysis "dynamic" even if the game has only one stage). Each α_i is such that

- $\forall h \in H$, $\forall z \in Z$, $\alpha_i(z|h) > 0$ only if $h \prec z$; $\forall z \in Z$, $\alpha_i(z'|\mathcal{P}_i(z)) > 0$ only if $z' \in \mathcal{P}_i(z)$ (i believes what she observes);
- *Chain rule:* α_i satisfies the rules of conditional probabilities when applicable (that is, if she did not assign prob. 0 to what she later observed); thus,

$$\alpha_i\left(\left(h, a', a''\right)|h\right) = \alpha_i\left(\left(h, a', a''\right) \mid \left(h, a'\right)\right)\alpha_i\left(\left(h, a'\right)|h\right);$$

- *Self* vs *others independence:* what i believes about others does not depend on her chosen actions; thus, $\alpha_i(h, (a_i, a_{-i})|h) = \alpha_{i,i}(a_i|h) \times \alpha_{i,-i}(a_{-i}|h)$, where $(\alpha_{i,i}(\cdot|h))_{h \in H_i}$ and $(\alpha_{i,-i}(\cdot|h))_{h \in H_{-i}}$ are the marg. prob. of actions given α_i.

Second-order beliefs

- We let Δ_i^1 denote the space of 1^{st}-order belief systems of $i \in I$. Second-order beliefs matter if players care about the 1^{st}-order beliefs of others: $u_i : Z \times \left(\times_{j\in I} \Delta_j^1 \right) \to \mathbb{R}$.
- Given each $h \in H \cup \mathcal{P}_i$, (real) pl. i has (conditional) 2^{nd}-order belief $\beta_i(\cdot|h) \in \Delta\left(Z \times \left(\times_{j\neq i} \Delta_j^1 \right) \right)$. 2^{nd}-order belief systems $\beta_i = (\beta_i(\cdot|h))_{h\in H\cup\mathcal{P}_i}$ describe i's initial and conditional beliefs about paths of play *and* the 1^{st}-order beliefs of others. They satisfy properties similar to those of 1^{st}-order belief systems.
- In particular, from $\beta_i = (\beta_i(\cdot|h))_{h\in H\cup\mathcal{P}_i}$ we derive a corresponding belief system $\alpha_i = (\alpha_i(\cdot|h))_{h\in H\cup\mathcal{P}_i}$ by marginalization: $\alpha_i(z|h) = \beta_i\left(\{z\} \times \left(\times_{j\neq i} \Delta_j^1 \right) |h \right)$. With this, β_i must be such that the derived α_i satisfies the aforementioned properties of 1^{st}-order beliefs, that is, $\alpha_i \in \Delta_i^1$.
- We let Δ_i^2 denote the space of 2^{nd}-order belief systems of i.

Best replies

- Best replies depend on information. Whenever active, a rational player chooses actions that are best replies to his *conditional* beliefs, which include how he planned to continue afterward: rational planning, or one-step optimality, or intrapersonal equilibrium.
- Formally, let $\bar{u}_{i,h}(a_i;\beta_i) = \mathbb{E}_{\beta_i}(u_i|h,a_i)$ denote the subjective expected utility of taking action a_i at any $h \in H_i$ given the conditional 2^{nd}-order belief $\beta_i(\cdot|h)$. **Rational planning** requires $\alpha_{i,i}(a_i|h) > 0 \Rightarrow a_i \in \arg\max_{a_i' \in A_i(h)} \bar{u}_{i,h}(a_i';\beta_i)$ for all $h \in H_i$ and $a_i \in A_i(h)$ (where α_i is derived from β_i).
- *If* u_i satisfies *own-plan independence* (that is, u_i does not depend on $\alpha_{i,i}$), *then* preferences over continuation plans given updated beliefs satisfy *dynamic consistency*. This implies, by a "folding-back" argument, that one-step optimality is equivalent to re-optimization over continuation plans starting from every h (*One-Deviation Principle*). We omit the details.

Exercises on rational planning

1. Find the rational plan $\alpha_{1,1}$ of pl. 1 in the *BoSOO* with *selfish and risk neutral* players in the following two cases: (i) $\alpha_{1,2}\left(\mathrm{B}_2|\mathrm{In}\right)=\frac{1}{2}$, (ii) $\alpha_{1,2}\left(\mathrm{B}_2|\mathrm{In}\right)=\frac{3}{4}$. Find the set of rational plans for (iii) $\alpha_{1,2}\left(\mathrm{B}_2|\mathrm{In}\right)=\frac{1}{4}$ and (iv) $\alpha_{1,2}\left(\mathrm{B}_2|\mathrm{In}\right)=\frac{2}{3}$.
2. Add to the *reporting game* form presented above the (parametric) utility function

$$u_1\left(x,y,\alpha_2\right)=y-\theta_1\sum_{x'=0}^{5}\alpha_2\left(x'|y\right)\left[y-x'\right]^{+}$$

(note, u_1 is independent of x, only 2's perception of cheating matters to pl. 1, see B&D 2020, Sec. 4.1). Suppose that pl. 1 is certain that report $y=0$ would be believed, and any $y>0$ would not, with all lower numbers deemed equally likely. Find the rational plan $\alpha_{1,1}$ for $\theta_1=1$ and $\theta_1=2$. Find the set of rational plans for $\theta=\frac{5}{3}$.

Predictions

- The dominant paradigm of traditional GT is to derive predictions using refinements (strengthenings) of the Nash equilibrium concept, such as "subgame perfect equilibrium", or "sequential equilibrium". Such methodology is of little use for predicting subjects' behavior in experimental settings.
- One can define notions of "iterated deletion of never (sequential) best replies". Sometimes 2-3 rounds of deletion are useful to derive meaningful predictions in experiments.

Predictions: a traditional-GT example

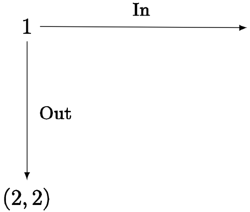

1\2	B_2	S_2
B_1	3, 1	0, 0
S_1	0, 0	1, 3

Suppose players in BoSOO are *commonly known to be selfish and risk neutral* (at least approximately, for the given stakes). Then

1. Delete plan $(\mathrm{In}, \mathrm{S}_1)$ (at most €1) dominated by Out (€2).
2. Delete plan S_2: indeed, if 2 maintains (whenever possible) the assumption that 1 is rational, then In "signals" that 1 will continue with B_1; the unique best reply is B_2.
3. Delete plan Out: if 1 reasons as above about 2, then his unique best reply is to play In with the plan to continue with B_1.

Predictions: a PGT example, I

- Consider the following Trust mini-Game form (G6 in B&D 2020): Ann (pl. 1) either **d**oesn't trust Bob (pl. 2), or **t**rusts Bob; in the latter case, the sum of payoffs doubles [from €(5+5) to €(10+10)] *if* equally **s**hared; but Bob can also **g**rab, reducing total payoff and keeping the rest (€14) for himself.

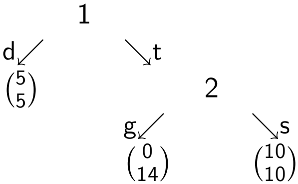

- Suppose p-utility functions have the (parametric) "guilt-averse" form $u_i(z, \alpha_j) = \pi_i(z) - \theta_i \left[\mathbb{E}_{\alpha_j}(\pi_j) - \pi_j(z)\right]^+$. Finally, suppose that it is commonly known between Ann and Bob (who know each other very well) that $\theta_1 = 0$ and $\theta_2 = 1$.

Predictions: a PGT example, II

It is common knowledge that

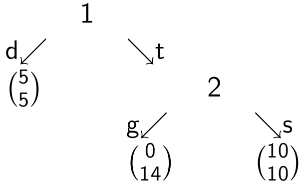

$$u_1 = \pi_1, \; u_2(\cdot, \alpha_1) = \pi_2 - [\mathbb{E}_{\alpha_1}(\pi_1) - \pi_1]^+$$

With this:

1. Ann (pl. 1) **t**rusts Bob (pl. 2) only if $\alpha_{1,2}(\mathrm{s}|\mathrm{t}) \geq \frac{1}{2}$; delete all (t, α_1) with $\alpha_{1,2}(\mathrm{s}|\mathrm{t}) < \frac{1}{2}$.
2. As Bob maintains (when possible) the assumption that Ann is rational, $\beta_2\left(\alpha_{1,2}(\mathrm{s}|\mathrm{t}) \geq \frac{1}{2}|\mathrm{t}\right) = 1$; delete **g**.
3. Ann understands this and **t**rusts Bob; delete **d**.

References

- Battigalli, P. (2020): *Game Theory: Analysis of Strategic Thinking*. Typescript, Bocconi University. [Downloadable from webpage, optional.]
- Battigalli, P. (2020): *Mathematical Language and Game Theory*. Typescript, Bocconi University. [Downloadable from webpage, optional.]
- Battigalli, P., C. Corrao, and M. M. Dufwenberg (2019): "Incorporating Belief-Dependent Motivation in Games," *Journal of Economic Behavior & Organization*, **167**, 185-218. [Downloadable from webpage, optional.]
- Battigalli, P., and M. Dufwenberg (2020): "Belief-Dependent Motivations and Psychological Game Theory," *Journal of Economic Literature*, forthcoming.