Backward Induction and Subgame Perfect Equilibrium

Pierpaolo Battigalli
Bocconi University
Game Theory: Analysis of Strategic Thinking

13 April 2026

Abstract

We define and illustrate the subgame perfect equilibrium (SPE) concept. We then relate it to “continuation rationalizability” and present algorithms to compute the set of SPEs of finite games, or at least games with finite horizon. We start with the “backward-induction” (BI) solution of leader-follower games, which is very simple and coincides with all versions of rationalizability for sequential games. Next we extend BI to finite games with perfect information and no relevant ties, where it finds the unique continuation rationalizable profile and SPE. BI is then compared with strong rationalizability. It turns out that, in the aforementioned games, strong rationalizability (like iterated admissibility) yields the BI path, not necessarily the BI strategies. Finally, we define an algorithm to find all the SPEs of two-stage games. [These slides summarize and, in part, complement Chapters 11.4 and 12 of GT-AST.]

Subgame Perfect Equilibrium and Backward Induction in Leader-Follower Games

- **Sequential duopoly:** Each firm can only produce a high or low output, but the firm of Ann moves first (leader) and the firm of Bob follows (follower):

Ann\Bob	low (Left)	high (Right)
high (Up)	3, 1	0, 0
low (Down)	2, 2	1, 3

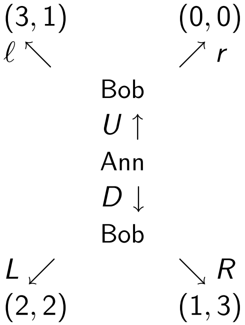

- $(D, r.R)$ is a Nash equilibrium. But Bob's rationality requires ℓ if U. **Traditional GT**: r *not an NE* of (U)-*subgame*⇒reject $(D, r.R)$!

Subgame Perfect Equilibrium

- A conjecture $\beta^i \in \times_{h\in H}\Delta\left(\mathcal{A}_{-i}(h)\right)$ is **degenerate** and corresponds to s_{-i} if $\beta^i\left(s_{-i}(h)\,|h\right) = 1$ for all $h \in H$.

Definition

A strategy profile s^* is a **subgame-perfect equilibrium** (SPE) if, for every $i \in I$, s_i^* is sequentially optimal given the degenerate conjecture corresponding to s_{-i}^*. [Using the seq-BR correspondence: $\forall i$, $s_i^* \in \hat{r}_i\left(s_{-i}^*\right)$.]

- A *similar definition* will be given for *randomized SPE* (see below).
- The OD Principle allows to characterize SPE by means of One-Step Optimality and to compute the SPE set in finite (horizon) games by means of a kind of backward-calculation procedure.
- Perfect information plays a special role. Recall, a game Γ has **perfect information (PI)** if, for each $h \in H$, *only one player*—denoted $\iota(h)$—*is active* at h: $\forall h \in H, \exists! i \in I, |\mathcal{A}_i(h)| > 1$.
- In such games, we represent *histories* as *sequences of actions*.

Backward Induction in Leader-Follower games

- **Definition.** A multistage game with observed actions Γ is a **leader-follower (LF) game** if it has
  - (i) *perfect information*,
  - (ii) *two stages* ($L(\Gamma) = 2$),
  - (iii) *two players*, and
  - (iv) the player who is active in the second stage (**follower**, by convention $i = 2$) is different from the player who is active in the first stage (**leader**, by convention $i = 1$).
- **Note:** In an LF game, actions and strategies of pl. 1 coincide: $S_1 = \mathcal{A}_1(\varnothing)$. Also, (i)-(ii) are essential, (iii)-(iv) are just simplifications.
- **Backward induction:**
  - Suppose that, in LF game Γ, for each $a_1 \in \mathcal{A}_1(\varnothing) \backslash Z$, the follower has a *unique best reply*, denoted $s_2^*(a_1) := \arg\max_{a_2 \in \mathcal{A}_2(a_1)} u_2(a_1, a_2)$.
  - Then, for every $s_1^* \in \arg\max_{a_1 \in \mathcal{A}_1(\varnothing)} u_1(a_1, s_2^*(a_1))$, $s^* = (s_1^*, s_2^*)$ is a SPE; also, *SPE and* (all versions of) *rationalizability coincide.*

Premise: Continuation Rationalizability

- Initial and strong rationalizability rely on the assumption that co-players' observed behavior—even when "surprising"—is interpreted as evidence about co-players' plans.
- **Continuation rationalizability (CR)** [aka "backward rationalizability"] instead *allows for the possibility that unexpected co-players' behavior is due to "mistakes"* in carrying out their plans, and yet no further "mistakes" will happen in the continuation-game, so that co-players' *continuation strategies* will be *consistent with rationality and "strategic reasoning" in the continuation game*. [Chapter 11.4 of GT-AST gives the details.]
- Continuation [aka "backward"] rationalizability characterizes the behavioral implications of *Rationality and "Common Future Belief in Rationality"*.

  [Battigalli & De Vito (2021) offer an in-depth analysis and foundation.]

Two-Stage Games

- In *two-stage games*, find the continuation rationalizable strategies with the following **backward procedure** [which explains the alternative "backward r." name]:
  - 1. First, find the *rationalizable actions* in each *last-stage* (second-stage) subgame.
  - 2. Next *iteratively delete* strategies of the two-stage game that are not best replies to conjectures assigning probability 0 to already deleted strategies, *taking into account Step 1*, i.e., co-players' strategies that select non-rationalizable actions in the second (last) stage must have probability 0.
- **Note**: The *three versions of rationalizability for multistage games coincide in* the even more special case of *leader-follower games* (two-stage games with perfect information).

Example

To ease reading, payoffs of player **1** are written in **bold**.

1

$\ell \swarrow$ $\searrow r$

Γ :

1\2	C	D
a	**3**, 1	**0**, 0
b	**0**, 1	**1**, 0

1\2	w	m	e
n	**0**, 0	**2**, 4	**1**, 1
c	**4**, 2	**0**, 0	**1**, 1
s	**1**,0	**1**, 0	**0**, 5

Applying rationalizability to (simultaneous-move) subgames $\Gamma(\ell)$ and $\Gamma(r)$ we obtain a reduced game:

1

ℓ ↙ ↘ r

1\2	C
a	**3**, 1

1\2	w	m
n	**0**, 0	**2**, 4
c	**4**, 2	**0**, 0

According to CR, *unlike SR*, observing r is *not* evidence that **1** will continue with c even if **1** could "secure" 3 utils with ℓ: r might be a mistake in trying to carry out strategy $\ell.a.n$, e.g., because **1** is certain of $C.m$.

With this insight find the CR strategies of the whole game Γ.

Backward Induction: Preliminaries

- **Backward induction** is an **algorithm** to compute continuation rationalizable (CR) strategies and SPEs in finite *perfect-information* games whenever the CR strategies and SPE are unique, hence in those with "no relevant ties":
- Given distinct terminal histories $z' \neq z''$, write $\pi(z', z'')$ for the player who is active at the longest common prefix (last common predecessor) of z' and z'' and, hence, is "**pivotal**" for reaching z' *vs* z''.
- A game with *perfect information* Γ has **no relevant ties (NRT)** if the active, pivotal player is never indifferent between distinct continuation paths:

$$\forall z', z'' \in Z,\ z' \neq z'' \Rightarrow u_{\pi(z',z'')}(z') \neq u_{\pi(z',z'')}(z'').$$

- **Note:** Many infinite, compact-continuous PI games have relevant ties (e.g., bargaining games). Hence, we first focus on *finite PI games with NRT*.

Backward Induction: Algorithm

Fix a *finite PI game with no relevant ties*. For each $h \in H$, we compute the (BR and) SPE-choice $s^*_{\iota(h)}(h)$ and the (BR and) SPE-value $V^*_j(h)$ for each j of h by induction on $L(\Gamma(h))$. Of course, we define $V^*_j(z) := u_j(z)$ for all $j \in I$ and $z \in Z$. With this:

- If $L(\Gamma(h)) = 1$, $\left(h, a_{\iota(h)}\right) \in Z$ for every $a_{\iota(h)} \in \mathcal{A}_{\iota(h)}(h)$; thus,
  - $i = \iota(h) \Rightarrow s^*_i(h) = \arg\max_{a_i \in \mathcal{A}_i(h)} V^*_i(h, a_i) = \arg\max_{a_i \in \mathcal{A}_i(h)} u_i(h, a_i)$ (unique by NRT),
  - $\forall j \in I$, $V^*_j(h) = V^*_j\left(h, s^*_{\iota(h)}(h)\right)$.
- Let $s^*_{\iota(h)}(h)$ and $V^*_j(h)$ be defined for every $h \in \bar{H}$ s.t. $L(\Gamma(h)) \leq k$ and every $j \in I$. If $L(\Gamma(h)) = k + 1$, then $L\left(\Gamma\left(h, a_{\iota(h)}\right)\right) \leq k$ for every $a_{\iota(h)} \in \mathcal{A}_{\iota(h)}(h)$; thus,
  - $i = \iota(h) \Rightarrow s^*_i(h) = \arg\max_{a_i \in \mathcal{A}_i(h)} V^*_i(h, a_i)$ (unique by NRT),
  - $\forall j \in I$, $V^*_j(h) = V^*_j\left(h, s^*_{\iota(h)}(h)\right)$.
- *By the OD principle, such s^* is the unique SPE of Γ.* (Alternatively, unique CR profile, and $SPE \subseteq CR$).

Example: Take It or Leave It

Budget of K€. Start from 0€. In each stage, 1€ is added to the pot; players alternate; the active player can take everything or leave it. At stage K s/he leaves to the other player (verify NRT).

$$(K=3)\qquad \begin{array}{ccccccc} \mathrm{a} & \xrightarrow{L_1} & \mathrm{b} & \xrightarrow{L_2} & \mathrm{a} & \xrightarrow{L_3} & \binom{0}{3} \\ \downarrow T_1 & & \downarrow T_2 & & \downarrow T_3 & & \\ \binom{1}{0} & & \binom{0}{2} & & \binom{3}{0} & & \end{array}$$

1. $V_{\mathrm{a}}^{*}(L_1,L_2)=\max\{0,3\}=3\ (T_3)$, $V_{\mathrm{b}}^{*}(L_1,L_2)=0$
2. $V_{\mathrm{b}}^{*}(L_1)=\max\{2,V_{\mathrm{b}}^{*}(L_1,L_2)\}=2\ (T_2)$, $V_{\mathrm{a}}^{*}(L_1)=0$
3. $V_{\mathrm{a}}^{*}(\varnothing)=\max\{1,V_{\mathrm{a}}^{*}(L_1)\}=1\ (T_1)$, $V_{\mathrm{b}}^{*}(\varnothing)=0$

- $s^*=(T_1.T_3,T_2)$ BI solution, unique SPE.
- **Note**: Strong rationalizability (like IA=iterated admissibility) deletes (1) $L_1.L_3$, (2) L_2, (3) $L_1.T_3$, and yields $\{T_1.T_3, T_1.L_3\}\times\{T_2\}$ ($\{T_1.T_3, T_1.L_3\}=\mathbf{T}_1$ is a reduced strategy, by *weak* seq. optimality).

Backward Induction and Rationalizability

- BI in (finite) PI-games with NRT has a "flavor" of rationality and common belief in rationality. Indeed, it yields the unique continuation rationalizable profile. Yet, in some games, BI differs in essential ways from both initial and strong rationalizability.
- **Example** (verify NRT):

$$\begin{array}{ccccccccc}
\mathrm{a} & \overset{C_1}{\longrightarrow} & \mathrm{b} & \overset{C_2}{\longrightarrow} & \mathrm{a} & \overset{C_3}{\longrightarrow} & \mathrm{b} & \overset{C_4}{\longrightarrow} & \binom{4}{0} \\
\downarrow D_1 & & \downarrow D_2 & & \downarrow D_3 & & \downarrow D_4 & & \\
\binom{3}{0} & & \binom{0}{1} & & \binom{1}{0} & & \binom{0}{2} & &
\end{array}$$

  - The *BI* (and CR) *solution* is $s^* = (D_1.D_3, D_2.D_4)$; $C_1.D_3$ and $C_2.C_4$ are conditionally dominated.
  - The *Initially Rationalizable set* is $\rho^\infty(S) = \rho^2(S) = \{\mathbf{D}_1\} \times \{\mathbf{D}_2, C_2.D_4\}$ (for $i = 1, 2$, $\mathbf{D}_i$ is a reduced strategy, $s_i^* \in \mathbf{D}_i$).
  - The *Strongly Rationalizable set* is $S^\infty = S^2 = \{\mathbf{D}_1\} \times \{C_2.D_4\}$. **Note**: $s^* \notin S^\infty$ (because $s_2^* \neq C_2.D_4$), but $\forall s \in S^\infty$, $\zeta(s) = \zeta(s^*)$.

Backward Induction and Strong Rationalizability

- Although BI and strong rationalizability (or IA) may select very different sets of (reduced) strategies, one can prove the following important result (a corollary of more general results, e.g., relating SR-paths and CR-paths):

Theorem

For every (finite) PI game with no relevant ties, every strongly rationalizable strategy profile induces the same terminal history as the unique SPE s^**, that is,* $\zeta(S^\infty) = \{\zeta(s^*)\}$.

- For example, in the ToL3 game $\zeta(S^\infty) = \{(T_1)\}$, and in the previous example $\zeta(S^\infty) = \{(D_1)\}$. In both cases, the unique BI-path is selected by strong rationalizability.
- The same holds with "strong rationalizability" replaced by "iterated admissibility".

Backward Induction and Folding Back: Discussion

- *Many scholars find BI very compelling. Yet*, BI does not represent the behavioral implications of compelling assumptions about sophisticated strategic reasoning in sequential games, such as Rationality and Common Strong Belief in Rationality.
  - Of course, most scholars are not exposed to foundational analysis studying versions of RCBR in sequential games. This plays a role.
  - But the *seeming compellingness of BI probably comes from its similarity to the Folding Back* algorithm. It is therefore *important to emphasize the conceptual difference between BI and FB.*
  - **FB** is an algorithm to find the *intra*-personal equilibrium of a sophisticated planner, who knows herself, hence knows what she would believe conditional on each $h \in H$, and is therefore able to predict how she would choose at each $h \in H$.
  - **BI** is an algorithm to find an *inter*-personal equilibrium among players who cannot know what others would think conditional on each h. One may think that such beliefs should be pinned down by sophisticated strategic thinking. But the only foundation of this kind, CR, rests on *questionable assumptions.*

"Case-by-Case" Backward Induction, 0/3: BoSOO

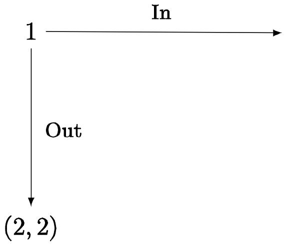

$1\backslash 2$	B_2	S_2
B_1	$3, 1$	$0, 0$
S_1	$0, 0$	$1, 3$

- How can we compute *all* the (pure) SPEs?
  - The BoS has 2 (pure) NEs: (B_1, B_2) and (S_1, S_2), BI-cases:
  - **Case** (B_1, B_2): pl. 1 thinks (B_1, B_2) would occur in BoS; then $V_1^{(B_1,B_2)}(\text{In}) = 3 > 2 = u_1(\text{Out}) \Rightarrow$ SPE $(\text{In}.B_1, B_2)$.
  - **Case** (S_1, S_2): pl. 1 thinks (S_1, S_2) would occur in BoS; then $V_1^{(S_1,S_2)}(\text{In}) = 1 < 2 = u_1(\text{Out}) \Rightarrow$ SPE $(\text{Out}.S_1, S_2)$.

"Case-by-Case" Backward Induction, 1/3

- BI can be extended to some games with multiple active players at some stages. E.g., for all the T-fold repetitions ($T < \infty$) of a static game G where each player i has a dominant action a_i^*, or a unique rationalizable action a_i^* (e.g., T-repeated PD), the CR/BI solution is s^* with $s_i^*(h) = a_i^*$ for every $h \in H$ and $i \in I$.
- *For games that do not have a unique SPE* computable by CR/BI, we still have a method to compute *all* the (pure) SPEs.
- We describe it for *two-stage games* (for simplicity, *finite*). Consider a *two-stage* game Γ. Fix any 1^{st}-stage, non-terminal (hence, pre-terminal) action profile $a^1 \in \mathcal{A}(\varnothing) \setminus Z$. Consider the **$2^{nd}$-stage subgame**

$$G^2(a^1) = \left\langle I, \left(\mathcal{A}_i(a^1), u_i(a^1, \cdot)\right)_{i \in I} \right\rangle$$

with payoff functions

$$\begin{array}{rccc} u_i(a^1, \cdot) : & \mathcal{A}(a^1) & \to & \mathbb{R} \\ & a^2 & \mapsto & u_i(a^1, a^2) \end{array}$$

"Case-by-Case" Backward Induction, 2/3

Stage-2 analysis: Let $NE^2\left(a^1\right)$:=set of (pure) Nash equilibria of $G^2\left(a^1\right)$.

- If $NE^2\left(a^1\right) = \emptyset$ for some $a^1 \in \mathcal{A}(\varnothing) \setminus Z$, then Γ cannot have any (pure) SPE.
- Suppose: $\forall a^1 \in \mathcal{A}(\varnothing) \setminus Z$, $NE^2\left(a^1\right) \neq \emptyset$. Consider all possible **selections** s^2 from the 2^{nd}-stage equilibrium correspondence $a^1 \mapsto NE^2\left(a^1\right)$, that is, all

$$s^2 \in \times_{a^1 \in \mathcal{A}(\varnothing)\setminus Z} NE^2\left(a^1\right)$$

(note: there are $\prod_{a^1 \in \mathcal{A}(\varnothing)\setminus Z} \left| NE^2\left(a^1\right) \right|$ such selections s^2).

"Case-by-Case" Backward Induction, 3/3

- **Stage-1 analysis**
  - Each $s^2 \in \times_{a^1 \in \mathcal{A}(\varnothing)\setminus Z} NE^2\left(a^1\right)$ is a "**case**" to which we apply a "backward induction" calculation, that is:
  - define the *auxiliary* simultaneous-move *game*

$$G^1\left(s^2\right) = \left\langle I, \left(\mathcal{A}_i\left(\varnothing\right), u_i^1\left(\cdot, s^2\right)\right)_{i \in I}\right\rangle,$$

    where

$$u_i^1\left(a^1, s^2\right) = \begin{cases} u_i\left(a^1, s^2\left(a^1\right)\right), & \text{if } a^1 \in \mathcal{A}\left(\varnothing\right)\setminus Z, \\ u_i\left(a^1\right), & \text{if } a^1 \in Z. \end{cases}$$

  - $G^1\left(s^2\right)$ specifies the payoffs of each first-stage action profile a^1 under the hypothesis (commonly believed by the players) that, for each $a^1 \in \mathcal{A}\left(\varnothing\right)\setminus Z$, the following 2^{nd}-stage profile will be $s^2\left(a^1\right)$.
  - Every NE $s^1 \in \mathcal{A}\left(\varnothing\right)$ of $G^1\left(s^2\right)$ yields a SPE $s = \left(s^1, s^2\right)$: $\forall i \in I$, $s_i\left(\varnothing\right) = s_i^1$, $\forall a^1 \in \mathcal{A}\left(\varnothing\right)\setminus Z$, $s_i\left(a^1\right) = s_i^2\left(a^1\right)$.
- The number of SPEs of Γ, $\left|SPE\left(\Gamma\right)\right|$, is the sum over "cases" s^2 of the numbers of equilibria in the auxiliary games $G^1\left(s^2\right)$.

Example

In game Γ, every a^1 except (D,R) is terminal, (D,R) leads to G^2:

a\b	L	R
U	2, 1	1, 2
D	1, 2	G^2

with

G^2	ℓ	r
u	**2, 1**	0, 0
d	0, 0	**1/2, 2**

.

Figure 2 Game Γ.

▶Two auxiliary games according to selected equilibrium of G^2:

$G^1(u,\ell)$	L	R
U	2, 1	1, 2
D	1, 2	2, 1

$G^1(d,r)$	L	R
U	2, 1	**1, 2**
D	1, 2	1/2, 2

.

▶ $G^1(u,\ell)$ has no equilibrium (like "Matching Pennies")⇒ no SPE where (u,ℓ) is selected in G^2. ▶ $G^1(d,r)$ is dominance-solvable⇒ unique eq. (U,R).▶ *Unique* SPE, $(U.d, R.r)$: Ann does not deviate to D in the 1^{st} stage because she expects to be "punished" by (d,r). ▲

Exercise

- Find the SPEs of the Battle of the Sexes with Dissipative Action.
  - What are the "cases" s^2 to be considered?
  - Compare the results with initial rationalizability, strong rationalizability (or iterated admissibility), and continuation rationalizability.

Randomized Subgame Perfect Equilibria

- The previous method can be extended to find **randomized SPEs**, i.e., subgame perfect equilibria in behavior strategies.

Definition

A profile of behavior strategies $\beta = (\beta_i)_{i \in I}$ is a **subgame perfect equilibrium** (**SPE**) if, for each $i \in I$, β_i is sequentially optimal given the (correct) conjecture $\beta^i = \beta_{-i}$.

- By the OD principle,
  - β is a *subgame perfect equilibrium IFF*

$$\forall h \in H, \forall i \in I, \mathrm{supp}\beta_i(\cdot|h) \subseteq \arg\max_{a_i \in \mathcal{A}_i(h)} V_i^{\beta}(h, a_i),$$

  - where $V_i^{\beta}(h, a_i) = \sum_{a_{-i} \in \mathcal{A}_{-i}(h)} \mathbb{P}^{\beta}(z|h, (a_i, a_{-i})) u_i(z)$.

Theorem

Every finite Γ *has a SPE in behavior strategies.*

Randomized Equilibria in 2-Stage Finite Games

- **Stage 2 analysis.** As in the previous algorithm, start finding each mixed equilibrium of each second-stage subgame $G^2(a^1)$ $(a^1 \in \mathcal{A}(\varnothing) \setminus Z)$. Let $MNE^2(a^1)$ denote the set of mixed Nash equilibria of $G^2(a^1)$ (for "almost" all finite games $\left|MNE^2(a^1)\right|$ is an odd number). Each $\beta^2 \in \times_{a^1 \in \mathcal{A}(\varnothing)\setminus Z} MNE^2(a^1)$ is a "**case**" to start from (here, β^2=stage-2 profile).
- **Stage 1 analysis.** For each "case" β^2 define $G^1(\beta^2) = \left\langle I, \left(\mathcal{A}_i(\varnothing), u_i^1(\cdot, \beta^2)\right)_{i \in I} \right\rangle$, where

$$u_i^1(a^1, \beta^2) = \begin{cases} \sum_{a^2 \in \mathcal{A}(a^1)} u_i(a^1, a^2) \prod_{j \in I} \beta_j^2\left(a_j^2 | a^1\right), & \text{if } a^1 \notin Z, \\ u_i(a^1), & \text{if } a^1 \in Z. \end{cases}$$

Find all the MNEs $\beta^1 \in \times_{i \in I} \Delta(\mathcal{A}_i(\varnothing))$ of $G^1(\beta^2)$. The profile β such that $\beta_i(\cdot|\varnothing) = \beta_i^1$ and $\beta_i(\cdot|a^1) = \beta_i^2(\cdot|a^1)$, for every $i \in I$ and $a^1 \in \mathcal{A}(\varnothing) \setminus Z$, is a SPE (here, β^1=stage-1 profile).

Exercises

- Find the (pure and) randomized SPEs of the BoS with an Outside Option.
- Find the (pure and) randomized SPEs of the BoS with a Dissipative Action.
- Find the (pure and) randomized SPEs of the previous example.

References

BATTIGALLI, P., E. CATONINI, AND N. DE VITO (2025): *Game Theory: Analysis of Strategic Thinking*. Typescript, Bocconi University.

BATTIGALLI, P. (2025): *Mathematical Language and Game Theory*. Typescript, Bocconi University.

BATTIGALLI, P., AND N. DE VITO (2021): “Beliefs, Plans and Perceived Intentions in Dynamic Games,” *Journal of Economic Theory*, 195, 105283. [For in-depth study.]