

14. Extensive-Form Games and Equilibrium Refinements

Duarte Gonçalves

University College London

MRes Microeconomics

Normal-Form Games and Extensive-Form Games

Limitations of Normal-Form Representations

Normal-form representation of games: simple, useful, but lacks notion of time.

Some players may be able to observe opponents' choices before making their own.

Examples:

- Employers may know which courses students chose to take.
- Banks observe central bank's monetary policy before deciding on loans.
- Firms may observe their competitors' pricing decisions before making theirs.
- Employers and employees 1st sign contracts, 2nd employees decide how much effort to exert, and 3rd firms decide on bonuses/promotions.
- Firms make choices about which technologies to invest in prior to start producing.

Normal-Form Games and Extensive-Form Games

Information Matters

Not only that actions may be dynamic, but how dynamics interacts with information

Two competing firms set prices for the following day.

If neither can observe competitor's price in advance, then exact time price gets set is of no consequence.

But if firm learns its competitor's pricing decision in advance, then it can condition its own price on opponent's price.

Crucial to capture what players know when making decisions;
otherwise model predictions could be very much at odds with the data.

Need a different way to model games to account for the fact that:

- (1) strategic interaction unfolds over time, and
- (2) what players know when they make their choices matters.

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
5. Applications
6. Beliefs and Sequential Rationality
7. Weak Perfect Bayesian Equilibrium
8. Sequential Equilibrium
9. More

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
 - History-Based Definition
 - Game Trees
 - Strategies
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
5. Applications
6. Beliefs and Sequential Rationality
7. Weak Perfect Bayesian Equilibrium
8. Sequential Equilibrium
9. More

Definition

An **extensive-form game** is given by a tuple $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

- (1) I : set of *players*; nature or chance is represented by $0 \notin I$.
- (2) $\mathcal{A}$: overall set of *actions*.
- (3) H : set of *histories*.
- (4) $\mathcal{H} := \{\mathcal{H}_i\}_{i \in I \cup \{0\}}$, where $\mathcal{H}_i$ is player i 's *information sets* or *information partition*.
- (5) ρ : function that associates each of nature's information sets $H_0 \in \mathcal{H}_0$ with probability measure over feasible actions after any history $h \in H_0$, $\rho(H_0) \in \Delta(A(H_0))$.
- (6) $u := (u_i)_{i \in I}$, where each u_i represents player i 's payoff function, $u_i : T \rightarrow \mathbb{R}$.

Extensive-Form Games

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

I denotes a set of players; nature or chance is represented by $0 \notin I$

As before, I is the set of players.

Nature to represent randomness (whether or not $\text{nature} \in I$ is just convention).

E.g. Firms decide on investment decisions; with some prob. a pandemic will start.

$\mathcal{A}$ denotes the overall set of actions.

All the actions that some player or nature can take at some point.

Extensive-Form Games

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

H denotes the set of histories, satisfying the following properties

- (i) The empty history $\emptyset$ is a member of H
(the 'starting point' of the game).
- (ii) A nonempty history $h \in H$ consists of a (possibly infinite) sequence of actions,
 $h = (a^1, \dots, a^t) \in \mathcal{A}^t$ for some $t \in \mathbb{N} \cup \{\infty\}$.
(what has happened thus far)

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

H denotes the set of histories, satisfying the following properties

(iii) If, for $n \in \mathbb{N} \cup \{\infty\}$, $(a^\ell)_{\ell=1}^n \in H$, then, for any positive integer $m < n$, $(a^\ell)_{\ell=1}^m \in H$.

A (proper) subhistory h' of history $h = (a^1, \dots, a^t)$ is a sequence of actions

$h' = (a'^1, \dots, a'^s)$ such that $s \leq (<)t$ and $a^n = a'^n$ for $n = 1, \dots, s$

(if a given seq of n actions is a feasible history, then so are its subhistories).

(iv) If $(a^\ell)_{\ell=1}^\infty$ is such that, for every $n \in \mathbb{N}$, $(a^\ell)_{\ell=1}^n \in H$, then $(a^\ell)_{\ell=1}^\infty \in H$

(if all finite subhistories are feasible histories, then so is the history).

Extensive-Form Games

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where
I players; $\mathcal{A}$ actions; *H* histories; $\mathcal{H}_i$'s info sets; ρ nature's move; u_i payoffs

H denotes the set of histories

A history $h \in H$ is said to be a **terminal history** if (a) $(h, a) \notin H$ for any $a \in \mathcal{A}$;
or (b) it is an infinite sequence of actions.

The **set of terminal histories** is denoted by $T \subset H$. Terminal history $\equiv$ Outcome.

A history which is not terminal ($h \in H \setminus T$) is called a **nonterminal history**.

(Q: why don't we just do $H =$ all possible sequences of actions from $\mathcal{A}$?)

The **set of feasible actions** following nonterminal history h is defined as

$$A(h) := \{a \in \mathcal{A} \mid (h, a) \in H\}.$$

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

$\mathcal{H} := \{\mathcal{H}_i\}_{i \in I \cup \{0\}}$, where $\mathcal{H}_i$ denotes player i 's information sets or information partition (including nature), satisfying the following properties:

- (i) $H_i \in \mathcal{H}_i$ is an information set; consists of a subset of nonterminal histories, $H_i \subseteq H \setminus T$.
- (ii) The set of all players information sets (including nature) $\cup_{i \in I \cup \{0\}} \mathcal{H}_i$ determines a partition over the set of all nonterminal histories i.e.,
 - (a) any two information sets are disjoint ($\tilde{H} \cap \hat{H} = \emptyset, \forall \tilde{H}, \hat{H} \in \cup_{i \in I \cup \{0\}} \mathcal{H}_i$); and
 - (b) the union of all information sets of all players (including nature) corresponds to the set of nonterminal histories ($H \setminus T = \cup_{i \in I} \{\tilde{H}_i \in \mathcal{H}_i\}$).

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

$\mathcal{H} := \{\mathcal{H}_i\}_{i \in I \cup \{0\}}$, where $\mathcal{H}_i$ denotes player i 's information sets or information partition (including nature), satisfying the following properties:

- (i) $H_i \in \mathcal{H}_i$ is an information set; consists of a subset of nonterminal histories, $H_i \subseteq H \setminus T$.
- (ii) The set of all players information sets (including nature) $\cup_{i \in I \cup \{0\}} \mathcal{H}_i$ determines a partition over set of all nonterminal histories.
In general, nature's information sets are singletons, corresponding to a single history.
- (iii) For any two histories belonging to the same information set, $h, h' \in H_i \in \mathcal{H}_i$, the set of feasible actions is the same, $A(h) = A(h') =: A(H_i)$.

Extensive-Form Games

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

$\mathcal{H} := \{\mathcal{H}_i\}_{i \in I \cup \{0\}}$, where $\mathcal{H}_i$ denotes *player i 's information sets or information partition (including nature)*

Idea: $\mathcal{H}_i$ represents what player i knows.

Two histories are in same info set = player i cannot distinguish between them.

After sequence of actions $h = (a^1, a^2, \dots, a^t) \in H_i$, player i knows some history in H_i was played, but cannot observe which.

That is why player i has to choose the same action following all histories in the same info set $h \in H_i$.

When does each player move?: Player i moves following each history h that belongs to some information set $H_i \in \mathcal{H}_i$.

Extensive-Form Games

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

$\mathcal{H} := \{\mathcal{H}_i\}_{i \in I \cup \{0\}}$, where $\mathcal{H}_i$ denotes *player i 's information sets or information partition (including nature)*

$A(H_i) := A(h)$; denote set of feasible actions after any history in information set H_i .

If, following two different histories belonging to the same information set, player i had different actions available, then would be able to distinguish between the histories.

Extensive-Form Games

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

$\mathcal{H} := \{\mathcal{H}_i\}_{i \in I \cup \{0\}}$, where $\mathcal{H}_i$ denotes *player i 's information sets or information partition (including nature)*

Game is of (im)perfect information if (not) all information sets are singletons.

Game is of perfect recall if players don't forget (i) what they know nor (ii) which actions they take. Formally,

- (1) If $h \in H_i$, then for any proper subhistory h' of h , $h' \notin H_i$.
- (2) Let $h, h' \in H_i$, and take any $\tilde{h}, \tilde{h}' \in \tilde{H}_i$ that are subhistories of h and h' , resp., belonging to the same information set of player i .
Then $(\tilde{h}, a)$ is a subhistory of h if and only if $(\tilde{h}', a)$ is a subhistory of h' .
(Player must remember action taken at info set $\tilde{H}_i$.)

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

ρ : function that associates each of nature's info sets $H_0 \in \mathcal{H}_0$ with a prob. measure over set of feasible actions following any history $h \in H_0$, $\rho(H_0) \in \Delta(A(H_0))$.

Nature moves after any h that belongs to some information set $H_0 \in \mathcal{H}_0$.

ρ determines what nature does at each information set.

Definition

Extensive-form game: $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ where

I players; $\mathcal{A}$ actions; H histories; $\mathcal{H}_i$ i 's info sets; ρ nature's move; u_i payoffs

$u := (u_i)_{i \in I}$, where each u_i represents **player i 's payoff function**, $u_i : T \rightarrow \mathbb{R}$

Payoffs realise after terminal histories.

We will assume that u_i corresponds to a von-Neumann–Morgenstern utility function (Bernoulli index) representing preferences of player i over terminal histories.

Representing Extensive-Form Games

Game tree: nodes, edges, and information sets

Nodes: each node corresponds to a different history

Root: empty history, 'starting point' of the game

Terminal nodes: terminal histories; typically labeled with players' payoffs

Non-Terminal nodes: correspond to non-terminal histories; nodes/histories at which a player makes a choice

Only one player makes a choice at any given node/following any given history


Edges: correspond to different actions the player choosing may take; typically labeled with the name of the corresponding actions

Information sets: correspond to histories a given player is unable to distinguish between; typically labeled with the name of the player that is choosing/active
Represented by grouping of non-terminal nodes (circling them, dashed lines)
The same player choosing at any node/history in the same information set


Definition in MGW and notes: literal definition of a game tree representation of a finite extensive-form game. Equivalent for finite games.

History-based definition in Osborne & Rubinstein (1) more meaningful, (2) more versatile (easy to accommodate infinitely repeated games)

Perfect Information



Imperfect Information



Extensive-Form Games

Strategies in Extensive-form Games


At each $H_i \in \mathcal{H}_i$, player i 's feasible **actions** are $A(H_i)$.

(Pure) Strategy for player i : $s_i : \mathcal{H}_i \rightarrow \mathcal{A}$ such that $s_i(H_i) \in A(H_i)$


A (pure) strategy of player i specifies a full contingent plan: which feasible action player i chooses at each information set.

Think about it as delegating decision to a representative

What are the histories/info sets/strategies/actions?



What are the histories/info sets/strategies/actions?



Extensive-Form Games

Strategies in Extensive-form Games

At each $H_i \in \mathcal{H}_i$, player i 's feasible **actions** are $A(H_i)$.

(Pure) Strategy for player i : $s_i : \mathcal{H}_i \rightarrow \mathcal{A}$ such that $s_i(H_i) \in A(H_i)$

A (pure) strategy of player i specifies a full contingent plan: which feasible action player i chooses at each information set.

Think about it as delegating decision to a representative


In games without nature moves,

a pure strategy profile $(s_i)_{i \in I}$ induces a unique terminal history
(multiple pure strategy profiles may induce the same terminal history).

In general (with nature moves, randomness),

a pure strategy profile induces a distribution over terminal histories.

What are the histories/info sets/strategies/actions?



Strategies: $S_A = \{H, L\}$, $S_B = \{(h|H, h|L), (l|H, h|L), (l|H, l|L), (l|H, l|L)\}$

Both $(H, (h|H, h|L))$ and $(H, (h|H, l|L))$ induce terminal history Hh

Extensive-Form Games

Strategies in Extensive-form Games

Mixed Strategy for player i : distribution over pure strategies, $\sigma_i \in \Delta(S_i) =: \Sigma_i$.

Behavioural Strategy for player i : distribution over actions at each information set,

$\lambda_i : \mathcal{H}_i \rightarrow \Delta(\mathcal{A})$ such that $\lambda_i(H_i)(a) = 0 \forall a \notin A(H_i)$.

(Can only randomise over strategies that are feasible at H_i).

Both mixed and behavioural strategies induce distributions over terminal histories.


Terminology

Degenerate mixed strategy: $\exists s_i : \sigma_i(s_i) = 1$; mixed strategies subsume pure.

Non-Degenerate mixed strategy: $\nexists s_i : \sigma_i(s_i) = 1$; same as pure strategy; mixed subsume pure.

Fully mixed strategy: $\forall s_i : \sigma_i(s_i) > 0$.

What are the histories/info sets/strategies/actions?



Strategies: $S_A = \{H, L\}$, $S_B = \{(h|H, h|L), (l|H, h|L), (h|H, l|L), (l|H, l|L)\}$.

$\sigma_A(H) = 1/3$, $\sigma_B((h|H, h|L)) = 1/4$, $\sigma_B((l|H, h|L)) = 3/4$.

$$P_\sigma(Hh) = \sigma_A(H)(\sigma_B((h|H, h|L)) + \sigma_B((h|H, l|L))) = 1/3(1/4 + 0) = 1/12.$$

$$P_\sigma(Hl) = \sigma_A(H)(\sigma_B((l|H, h|L)) + \sigma_B((l|H, l|L))) = 1/3(0 + 3/4) = 3/12.$$

$$P_\sigma(Lh) = \sigma_A(L)(\sigma_B((h|H, h|L)) + \sigma_B((l|H, h|L))) = 2/3(1/4 + 3/4) = 2/3.$$

$$P_\sigma(Ll) = \sigma_A(L)(\sigma_B((h|H, l|L)) + \sigma_B((l|H, l|L))) = 2/3(0 + 0) = 0.$$

$\lambda_A(\emptyset)(H) = 1/3$, $\lambda_B(H)(h) = 1/4$, $\lambda_B(L)(h) = 1$.

Extensive-Form Games

Strategies in Extensive-form Games

Mixed Strategy for player i : distribution over pure strategies, $\sigma_i \in \Delta(S_i)$.

Behavioural Strategy for player i : distribution over actions at each information set,

$\lambda_i : \mathcal{H}_i \rightarrow \Delta(\mathcal{A})$ such that $\lambda_i(H_i)(a) = 0 \forall a \notin A(H_i)$.

(Can only randomise over strategies that are feasible at H_i .)

Both mixed and behavioural strategies induce distributions over terminal histories

Theorem (Kuhn's Theorem)

For finite extensive-form games with perfect recall, every mixed strategy of a player has an outcome-equivalent behavioural strategy and vice-versa.

To an extent, can use mixed and behavioural strategies interchangeably.

Note: What OR call Kuhn's theorem (Prop 99.2) is known as Zermelo's theorem (later).

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
5. Applications
6. Beliefs and Sequential Rationality
7. Weak Perfect Bayesian Equilibrium
8. Sequential Equilibrium
9. More

Nash Equilibria in Extensive-Form Games

A strategy profile $\sigma = (\sigma_i)_{i \in I}$ maps to a distribution over terminal histories $h \in T$.

Although $u_i : T \rightarrow \mathbb{R}$, we can unambiguously write $u_i : S \rightarrow \mathbb{R}$ (just as in normal-form games).

We also extend payoffs to mixed strategy profiles as before,

$$u_i(\sigma) := \sum_{s \in S} (\prod_{j \in I} \sigma_j(s_j)) u_i(s).$$

Definition

A **Nash equilibrium** of an extensive-form game $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ is a strategy profile $\sigma \in \Sigma$ such that for every player $i \in I$


$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}) \quad \forall \sigma'_i \in \Sigma_i.$$

Leverage already known existence results:

Proposition

Every finite ($|H| < \infty$) extensive-form game has a Nash equilibrium, possibly in mixed strategies.

From Extensive- to Normal-form



		Player B			
		$h H, h L$	$l H, h L$	$h H, l L$	$l H, l L$
Player A	H	2,2	1,3	2,2	1,3
	L	3,1	3,1	0,0	0,0

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
 - Credibility
 - Subgames
 - Subgame-Perfect Nash Equilibrium and Backward Induction
5. Applications
6. Beliefs and Sequential Rationality
7. Weak Perfect Bayesian Equilibrium
8. Sequential Equilibrium
9. More

The Problem of Credibility




Figure: Entry Game in Extensive-form

		Incumbent	
		fight	accommodate
Entrant	enter	-3,-1	2,1
	x enter	0,2	0,2

Figure: Entry Game in Normal-form

The Problem of Credibility

		Incumbent	
		fight	accommodate
Entrant	enter	-3,-1	2,1
	x enter	0,2	0,2

Figure: Entry Game in Normal-form

PSNE: (x, f) and (e, a) .

NE? $(\sigma_E(e), \sigma_I(a)) \{(0, p), p \in [0, 3/5]\} \cup \{(1, 1)\}$.

The Problem of Credibility




Figure: Entry Game in Extensive-form

PSNE: (x, f) and (e, a) .

But... (x, f) supported by I 's threat of fighting if E enters

Is this really credible? Not really.

Conditional on the entrant having entered, the incumbent is strictly better off accommodating.

Foreseeing this, entrant would choose to enter.

Incumbent is threatening to play f , but threat not credible.

(Note importance of having specified a full contingent plan!)

Definition

A **subgame** of an extensive-form game $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ is another extensive-form game $\Gamma(h) = \langle I, \mathcal{A}, H^h, \mathcal{H}(h), \rho^h, u^h \rangle$ such that

- (i) $\exists H_i = \{h\} \in \mathcal{H}_i$ s.t. $H^h = \{h' \mid h \text{ is subhistory of } h' \in H\}$; (ii) $\mathcal{H}_i^h \subseteq \mathcal{H}_i \forall i \in I$;
- (iii) $\rho^h(H_0) = \rho(H_0)$ for all $H_0 \in \mathcal{H}_0$; and (iv) $u_i^h(h') = u_i(h')$ for all $h \in T^h$.

For simplicity, write $\Gamma(h)$ for subgame starting following history h .

- (i) states that we start a subgame starts at a singleton information set of the game and includes all histories 'starting from there'; this implies that $T^h = T \cap H^h$;
- (subgame includes all its successors and nothing more)
- (ii) implies subgames don't 'cut across' information sets (players know they are playing the subgame);
- (iii) says nature moves the same way in the subgame as in the original game; and
- (iv) means that payoffs over the subgame's terminal histories are the same as in the original game.

Definition

A **subgame** of an extensive-form game $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ is another extensive-form game $\Gamma(h) = \langle I, \mathcal{A}, H^h, \mathcal{H}(h), \rho^h, u^h \rangle$ such that


- (i) $\exists H_i = \{h\} \in \mathcal{H}_i$ s.t. $H^h = \{h' \mid h \text{ is subhistory of } h' \in H\}$; (ii) $\mathcal{H}_i^h \subseteq \mathcal{H}_i \forall i \in I$;
- (iii) $\rho^h(H_0) = \rho(H_0)$ for all $H_0 \in \mathcal{H}_0$; and (iv) $u_i^h(h') = u_i(h')$ for all $h \in T^h$.

Remark


Let Γ be an extensive-form game.

- (1) $\Gamma = \Gamma(\emptyset)$.
- (2) Γ is a subgame of itself.
- (3) Any subgame of Γ is an extensive-form game.
- (4) If Γ is finite ($|H| < \infty$), then it has finitely many subgames.


How Many Subgames?




How Many Subgames?




How Many Subgames?



How Many Subgames?



How Many Subgames?



Definition

A **subgame** of an extensive-form game $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \rho, u \rangle$ is another extensive-form game $\Gamma(h) = \langle I, \mathcal{A}, H^h, \mathcal{H}(h), \rho^h, u^h \rangle$ such that

- (i) $\exists H_i = \{h\} \in \mathcal{H}_i$ s.t. $H^h = \{h' \mid h \text{ is subhistory of } h' \in H\}$; (ii) $\mathcal{H}_i^h \subseteq \mathcal{H}_i \forall i \in I$;
- (iii) $\rho^h(H_0) = \rho(H_0)$ for all $H_0 \in \mathcal{H}_0$; and (iv) $u_i^h(h') = u_i(h')$ for all $h \in T^h$.

Remark

Let G be the set of all subgames of Γ and let $\geq_g \subseteq G^2$:

$\Gamma(h) \geq_g \Gamma(h')$ iff $\Gamma(h')$ is a subgame of $\Gamma(h)$.

(1) G is nonempty for any Γ .

(2) $\geq_g$ is a partial order (reflexive, transitive, antisymmetric).

(Is $(G, \geq_g)$ a lattice?)

(2) $\forall \Gamma(h), \Gamma(h') \in G : \Gamma(h) \vee_g \Gamma(h')$ exists and is in G .

Also: $\Gamma(h) \wedge_g \Gamma(h')$ exists iff $\Gamma(h) \geq_g \Gamma(h')$ or vice-versa.

$(G, \geq_g)$ is not a lattice.


Refining Nash Equilibria in Extensive-Form Games: Subgame Perfection

Definition

A **subgame-perfect Nash equilibrium** (SPNE) of an extensive-form game Γ is a strategy profile σ that induces a Nash equilibrium in every subgame of Γ .

The whole game is a subgame of itself $\implies$ an SPNE is an NE.

Backward Induction



Backward Induction:

Main gist: start with terminal nodes/histories, pick payoff-maximising actions, and work your way backward

PSNE: $\{(H, (l|H, l|L)), (L, (h|H, h|L)), (L, (l|H, h|L))\}$

PSNE of subgame starting at H : l ; PSNE of subgame starting at L : h

PS-SPNE: $\{(L, (l|H, h|L))\}$

Definition (Backward Induction)

Let $G_0 := \emptyset$, $G^0 := \emptyset$, and $G_1 := \arg \min_{\geq_g} G \setminus G^0$.

$\forall k \in \mathbb{N}$

- Define $G^k := G^{k-1} \cup G_{k-1}$ and $G_k := \arg \min_{\geq_g} G \setminus G^k$.
- $\forall \Gamma' \in G_k$, pick a PSNE s' such that its implied behavioural strategies λ' are, at any information set, consistent with those fixed at any $\Gamma'' \in G_\ell \forall \ell < k$ where Γ'' is a subgame of Γ' .

s is obtained by backward induction if it results from the above procedure.


For any finite extensive-form game of perfect information, there is a strategy profile obtained by generalised backward induction.

Theorem (Zermelo's Theorem)

Let Γ be a finite ($|H| < \infty$) extensive-form game of perfect information.

- (1) Any s obtained by backward induction is a PSNE.
- (2) $\exists$ PSNE s that can be obtained by backward induction.
- (3) If no player has the same payoffs at any two terminal histories, then backward induction results in a unique strategy profile s .

Generalised Backward Induction



Generalised Backward Induction:

Main gist: start with subgames 'closest' to terminal nodes/histories, pick a NE in the subgame, and work your way backward

In this case, (u, r) is the unique NE of the only proper subgame

Definition (Generalised Backward Induction)

Let $G_0 := \emptyset$, $G^0 := \emptyset$, and $G_1 := \arg \min_{\geq_g} G \setminus G^0$.

$\forall k \in \mathbb{N}$

- Define $G^k := G^{k-1} \cup G_{k-1}$ and $G_k := \arg \min_{\geq_g} G \setminus G^k$.
- $\forall \Gamma' \in G_k$, pick a PSNE σ' such that its implied behavioural strategies λ' are, at any information set, consistent with those fixed at any $\Gamma'' \in G_\ell \forall \ell < k$ where Γ'' is a subgame of Γ' .

σ is obtained by backward induction if it results from the above procedure.

For any finite extensive-form game of perfect information, there is a strategy profile obtained by generalised backward induction.

Refining Nash Equilibria in Extensive-Form Games: Subgame Perfection

Definition

A **subgame-perfect Nash equilibrium** (SPNE) of an extensive-form game Γ is a strategy profile σ that induces a Nash equilibrium in every subgame of Γ .

Theorem

Let Γ be a finite ($|H| < \infty$) extensive-form game.

- (1) s is PS-SPNE if and only if it can be obtained by backward induction.
- (2) σ is SPNE if and only if it can be obtained by generalised backward induction.

Corollary

Let Γ be a finite ($|H| < \infty$) extensive-form game.

- (1) Γ has an SPNE.
- (2) If Γ is of perfect information, then Γ has PS-SPNE.
- (3) If Γ is of perfect information and no player has the same payoffs at any two terminal histories, then $\exists!$ SPNE. Furthermore, it is in pure strategies.

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
5. Applications
 - Alternating Bargaining
 - Centipede
6. Beliefs and Sequential Rationality
7. Weak Perfect Bayesian Equilibrium
8. Sequential Equilibrium
9. More

Alternating Bargaining

How to divide surplus?

Surplus generated by work, investment, trade, etc.

Nash (1950 Ecta) bargaining solution: outcome of axiomatic characterisation of desirable properties is simply to split it in half.

Can accommodate varying bargaining power, multiple players, etc.

Rubinstein (1982 Ecta): provide noncooperative foundation to Nash (1950 Ecta) bargaining solution.

Alternating Bargaining

Setup

- Two players, 1 and 2, bargain over the split of $\text{£}v > 0$.
- Up to T (odd) periods of bargaining.
- Both players discount payoffs at a rate $\delta \in (0, 1)$ per period.
- Conditional on bargaining continuing up to period t , Player i gets to propose a split $b_t \in [0, v]$, which the opponent can accept or reject, where $i = 1$ if t is odd and $i = 2$ if otherwise.
- If the opponent accepts, the game ends; proposer gets $\delta^{t-1}(v - b_t)$, and the opponent $\delta^{t-1}b_t$.
- If the opponent rejects, the game moves on to the next period $t + 1$ if $t < T$, or it ends if $t = T$, in which case both players get zero.
- Strategies are complicated, as they can depend on the whole observed history.

Alternating Bargaining

Backward Induction

- If $T = 1$, this is just an ultimatum game; same for $t = T$.
- At period T , Player 2 accepts if $b_T > 0$; if $b_T = 0$, Player 2 is indifferent.
- The unique SPNE in any subgame that reached period T is to have Player 1 proposing $b_T = 0$ and Player 2 accepting iff $b_T \geq 0$.
- $\forall b_T > 0$, Player 2 strictly prefers accepting over rejecting; hence Player 1 strictly prefers proposing $\frac{1}{2}b_T$ to get a higher share.
- Players accrue payoffs $\delta^{T-1}(v, 0)$.

Alternating Bargaining

Backward Induction

- At any subgame starting at period $T - 1$, Player 1 is willing to accept b_{T-1} iff

$$\delta^{T-2}b_{T-1} \geq \delta^{T-1}v \iff b_{T-1} \geq \delta v.$$

Otherwise, would prefer to reject and move to the next period and get the chance to propose.

- By similar argument, the unique SPNE in this subgame is to have Player 2 offering exactly

$$b_{T-1} = \delta v.$$

(This is but a sketch of the argument; requires a proper proof.)

- Payoffs are

$$(\delta^{T-1}v, \delta^{T-2}(1 - \delta)v).$$

Alternating Bargaining

Backward Induction

- At any subgame starting at period $T - 2$, the unique SPNE in the subgame has Player 1 proposing a split that Player 2 accepts while indifferent between accepting and rejecting.

$$\delta^{T-3}b_{T-2} = \delta^{T-2}(1 - \delta)v \iff b_{T-2} = \delta(1 - \delta)v.$$

- Payoffs are

$$\delta^{T-3}((1 - \delta + \delta^2)v, (\delta - \delta^2)v).$$

Alternating Bargaining

Backward Induction

- Iterating backward (via induction argument), at any $t \in [T - 1]$, the proposer suggests a split

$$b_{T-t} = \sum_{\ell=1}^t (-1)^{\ell-1} \delta^{\ell} v = \delta \frac{1 - (-1)^t \delta^t}{1 + \delta} v$$

and, opponent accepts iff

$$b_{T-t} \geq v \delta \frac{1 - (-1)^t \delta^t}{1 + \delta}.$$

- SPNE payoffs for the whole game are then

$$(v - b_1, b_1) = v \left(1 - \frac{\delta + (-1)^T \delta^T}{1 + \delta}, \frac{\delta + (-1)^T \delta^T}{1 + \delta} \right) = v \left(\frac{1 + \delta^T}{1 + \delta}, \frac{\delta - \delta^T}{1 + \delta} \right)$$

Alternating Bargaining

Insights

1. Unique SPNE!
2. No delay: a solution is reached immediately.
3. First and last propose confers advantage: Player 1 gets larger share of fixed resource.
As $T \rightarrow \infty$, equilibrium payoffs are given by $\left(v_{\frac{1}{1+\delta}}, v_{\frac{\delta}{1+\delta}}\right)$.
4. Patience pushes in favor of the last proposer; impatience, of the first.

Significance


Bread-and-butter of IO, macro-labour, etc.

Nash-in-Nash bargaining (when there are many parties negotiating at once); see Horn & Wolinsky (1988 RAND).

E.g., vertical integration health care market: Grennan 2013 AER; Ho & Lee 2017 Ecta.

Recent innovations: Noncooperation foundation by Collard-Wexler, Gowrisankaran, & Lee (2019 JPE).


Centipede Game



Setup

- Two players 1 and 2 take turns in choosing whether to continue or to stop
- Player 1 moves first; Player 2 moves after Player 1 provided Player 1 decided to continue, and vice-versa
- The game reaches a terminal node if either player decides to stop, or after each player decided to continue T times


Centipede Game



Setup. Payoffs are given as follows:

- Each player start with £1 in their pile
- Every time each player decides to continue, £1 is deducted from their pile and £2 are added to their opponents
- Their payoff equals the amount of money they have in their pile at the time they reach a terminal node

Centipede Game



Setup

$$A_i := \{0, 1\}, S_i := A_i^T, s_i = (a_{i,t})_{t \in [T]}$$

Writing payoffs formally is cumbersome... Let $a_{1,0} = a_{2,0} = 1$.

$$u_1(s_1, s_2) := 1 + 2 \sum_{t \in [T]} \left(\prod_{\ell \in [t]} a_{1,\ell} a_{2,\ell} \right) - \sum_{t \in [T]} \left(\prod_{\ell \in [t]} a_{1,\ell} a_{2,\ell-1} \right).$$

$$u_2(s_1, s_2) := 1 + 2 \sum_{t \in [T]} \left(\prod_{\ell \in [t]} a_{1,\ell} a_{2,\ell-1} \right) - \sum_{t \in [T]} \left(\prod_{\ell \in [t]} a_{1,\ell} a_{2,\ell} \right).$$



Centipede Game

Analysis

- The last subgame has Player 2 can either decide between continuing and getting $\pounds 1 + T - 1$ and stopping and getting $\pounds 1 + T$; then, by backward induction, $a_{2,T} = 0$.
- Then, as Player 2 stops in the last subgame, Player 1 prefers to stop and get $\pounds 1 + T - 1$, rather than continuing and get $\pounds 1 + T - 1 - 1$.
- Iterating backward, we'll find that the unique subgame perfect equilibrium has both players always stopping and getting $\pounds 1$!

Zermelo's theorem: no two terminal histories with the same payoff, hence unique SPNE, obtained by backward induction

Centipede Game



Experimental results (McKelvey and Palfrey, 1992)

Payoff frequencies (McKelvey & Palfrey 92 Ecta)

Centipede Game

Why?

- Monetary payoffs don't capture how players evaluate the outcome
This doesn't dent at the theory then: we just have the wrong payoff function.
- People may have limited foresight (inability to reason many steps ahead) and rely on heuristics.
Forward-looking behaviour often requires considering many contingencies, making issues fairly complicated.

Centipede Game




FIGURE 1. A CENTIPEDE GAME

(Palacios-Huerta & Volij 09 AER)

Centipede Game




FIGURE 2. COLLEGE STUDENTS: PROPORTION OF OBSERVATIONS AT EACH TERMINAL NODE

(Palacios-Huerta & Volij 09 AER)

Centipede Game




FIGURE 3. CHESS PLAYERS: PROPORTION OF OBSERVATIONS AT EACH TERMINAL NODE BY TYPE OF PLAYER 1 IN THE PAIR

Chess players stop earlier, and the earlier the higher their ranking.

(Palacios-Huerta & Volij 09 AER)

Centipede Game

TABLE 3—CHESS PLAYERS: IMPLIED STOP PROBABILITIES AT EACH TERMINAL NODE

	p_1	p_2	p_3	p_4	p_5	p_6	p_7
Grandmasters	1.00 (26)	1.00 (5)	—	—	—	—	—
International Masters	0.76 (29)	0.90 (10)	1.00 (2)	—	—	—	—
Federation Masters	0.73 (15)	0.66 (3)	1.00 (1)	—	—	—	—
Other chess players	0.61 (141)	0.58 (48)	0.73 (19)	0.80 (5)	1.00 (1)	—	—

Note: The number of players observed making a decision (stop or continue) at each node is in parentheses.

(Palacios-Huerta & Volij 09 AER)

Centipede Game

Why?

- Monetary payoffs don't capture how players evaluate the outcome
This doesn't dent at the theory then: we just have the wrong payoff function.
- People may have limited foresight (inability to reason many steps ahead) and rely on heuristics.
Forward-looking behaviour often requires considering many contingencies, making issues fairly complicated.
- How soon they stop depends on their beliefs on their opponent's strategic sophistication.
Stop later the less strategically sophisticated they perceive their opponent.

Centipede Game


TABLE 4—EXPERIMENTAL DESIGN FOR LABORATORY EXPERIMENT

Treatment	Subject pool Player 1 (white)	Subject pool Player 2 (black)	Session	Subjects	Games per subject	Total games
I	Students	Students	1	20	10	100
			2	20	10	100
II	Students	Chess players	3	20	10	100
			4	20	10	100
III	Chess players	Students	5	20	10	100
			6	20	10	100
IV	Chess players	Chess players	7	20	10	100
			8	20	10	100

(Palacios-Huerta & Volij 09 AER)

Centipede Game

Panel A. Students



Students learn to stop earlier *when playing with chess players*.

(Palacios-Huerta & Volij 09 AER)

Centipede Game




FIGURE 4. PERCENTAGE OF “STOP” IN NODE 1 PER ROUND

Chess players stop earlier *when playing other chess players* (and the earlier the higher their opponent’s ranking).


Also learn faster to stop earlier.

(Palacios-Huerta & Volij 09 AER)

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
5. Applications
6. Beliefs and Sequential Rationality
 - Credibility, Part 2
 - Belief System
 - Sequential Rationality
 - Restrictions on Beliefs
7. Weak Perfect Bayesian Equilibrium
8. Sequential Equilibrium
9. More

Credibility, Part 2



Only one subgame: the whole game! SPNE = NE.

(x, f) SPNE, but, in a sense, it's non-credible threat:

If incumbent is called to move, it would not be payoff maximising to choose to fight.

Want to have a way to rule out such equilibria.

Definition

A **belief system** (or system of beliefs) in an extensive-form game $\Gamma = \langle I, \mathcal{A}, H, \mathcal{H}, \mathbf{p}, u \rangle$ specifies, for each information set of each player, $H_i \in \mathcal{H}_i$, a probability distribution over the histories in that information set, $\mu(H_i) \in \Delta(H_i)$.

E.g., for $h \in H_i$, $\mu(H_i)(h)$ determines belief that player i holds upon being called to play at information set H_i that history h has occurred, conditional on information set H_i having been reached (i.e. player i having been called to play at information set H_i).

When $H_i = \{h\}$ contains only one history, $\mu(H_i)(h) = 1$
(no uncertainty on what happened before)

Sequential Rationality

Sequential rationality is a simple concept: each player, when called upon to play, chooses best behavioural strategy given their beliefs about what has happened and given what opponents are doing henceforth.

- $\mu(H_i)$: distribution over histories h in H_i .
- σ : induce distribution over terminal histories T .
- $T \upharpoonright_{H_i}$: terminal histories $h \in T$ s.t. $\exists$ history $h' \in H_i$ that is proper subhistory of h .
(i.e., terminal histories that follow from some history in H_i)
- $\sigma \upharpoonright_{H_i}$: distribution over terminal histories $T \upharpoonright_{H_i}$.
- $u_i \upharpoonright_{H_i}$: payoff function of player i restricted to $T \upharpoonright_{H_i}$.
- $\mathbb{E}[u_i(\sigma_i, \sigma_{-i}) \mid H_i, \mu]$: player i 's expected payoff at information set H_i given belief system μ and strategy profile σ .
(more properly, $\mathbb{E}[u_i \upharpoonright_{H_i}(\sigma \upharpoonright_{H_i}) \mid \mu(H_i)]$, but too cumbersome to carry around H_i)

Definition

Strategy profile σ is **sequentially rational at information set** H_i *given a belief system* μ if


$$\mathbb{E}[u_i(\sigma_i, \sigma_{-i}) \mid H_i, \mu] \geq \mathbb{E}[u_i(\sigma'_i, \sigma_{-i}) \mid H_i, \mu]$$

for all $\sigma'_i \in \Sigma_i$.

Strategy profile is **sequentially rational** *given a belief system* if it is sequentially rational at all information sets given that belief system.

Beliefs matter! Different μ can lead to different strategy sequentially rational strategies
It is sequentially rational *given* a belief system

Credibility, Part 2




(x, f) SPNE, but, in a sense, it's non-credible threat:

No belief that incumbent (P2) may hold at $H_2 = \{e_1, e_2\}$ would justify choosing to fight if called upon to move.

Given they are at H_2 , a is strictly better than f for *any beliefs* about whether the entrant (P1) chose e_1 ($\mu(H_2)(e_1)$) or e_2 ($\mu(H_2)(e_2) = 1 - \mu(H_2)(e_1)$).

Never sequentially rational at H_2 to choose f with positive probability.

Credibility, Part 2



Now (x, f) may be sequentially rational at H_2 given μ , but only if

$$\mathbb{E}[u_2(f, x) \mid H_2, \mu] \geq \mathbb{E}[u_2(a, x) \mid H_2, \mu] \iff p2 + (1 - p)(-1) \geq p0 + (1 - p)1 \iff p \in [1/2, 1]$$

where $p = \mu(H_2)(e_1) = 1 - \mu(H_2)(e_2)$.

Restrictions on Beliefs

For simplicity, focus on case where H is finite.

Recall σ determines prob. history h (a sequence of actions) is played.

Definition

An information set H_i is **reached given σ** if there is a positive probability that some history $h \in H_i$ is played with strictly positive probability, $\mathbb{P}(H_i \mid \sigma) > 0$

'Off-path': info set that is not reached according to σ .

'On-path': info set that is reached according to σ .

Definition

Belief system is **derived through Bayes rule whenever possible given σ** if, for any information set H_i that is reached given σ , beliefs $\mu(H_i)$ over histories in H_i equal the distribution over histories in H_i conditional on H_i , as induced by σ .

If H is finite, belief system is derived through Bayes rule whenever possible given σ if, whenever $\mathbb{P}(H_i \mid \sigma) > 0$,

$$\mu(H_i)(h) = \frac{\mathbb{P}(h \mid \sigma)}{\mathbb{P}(H_i \mid \sigma)}.$$


To be able to use Bayes rule, we need that, according to σ ,

[prob. history h being played given some history in H_i was played] is well-defined.

With finitely many histories, amounts to some $h \in H_i$ being played wp > 0 given σ

($\mathbb{P}(H_i \mid \sigma) > 0$), otherwise denominator on RHS = 0 and Bayes rule is not well-defined.

Credibility, Part 2




Suppose $\sigma_1(e_1) = 1/6$, $\sigma_1(e_2) = 1/3$, $\sigma_1(x) = 1/2$.

For $\mu(H_2)$ to be derived by Bayes' rule:

$$\mu(H_2)(e_1) = \frac{P(e_1 | \sigma)}{P(H_2 | \sigma)} = \frac{\sigma_1(e_1)}{\sigma_1(e_1) + \sigma_1(e_2)} = \frac{1/6}{1/6 + 1/3} = 1/3$$

and then $\mu(H_2)(e_2) = 1 - \mu(H_2)(e_1) = 2/3$.

Credibility, Part 2



Suppose $\sigma_1(e_1) = 0$, $\sigma_1(e_2) = 0$, $\sigma_1(x) = 1$.

Then, $\mu(H_2)$ cannot be derived by Bayes' rule from σ as $P(H_2 \mid \sigma) = 0$: H_2 is never reached given σ .

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
5. Applications
6. Beliefs and Sequential Rationality
7. Weak Perfect Bayesian Equilibrium
 - Comparison to Nash Equilibrium
 - Examples
 - Comparing WPBE and SPNE
 - Perfect Bayesian Equilibrium
8. Sequential Equilibrium
9. More

Refining Nash Equilibrium by Sequential Rationality

Definition

Strategy profile σ and belief system μ form a **weak perfect Bayesian Nash equilibrium** (WPBE) (σ, μ) of an extensive-form game Γ if

- (i) σ is sequential rational given the belief system μ ; and
- (ii) the belief system μ is derived through Bayes rule whenever possible given σ .

WPBE: need to define both the strategy profile **and the belief system**.

Beliefs are *required* to be correct *on-path*.

Care is needed in defining mixed strategies when there a given node has uncountably many successors; see Aumann (1964) 'Mixed and Extensive Strategies in Infinite Extensive Games'

Reinterpreting Nash Equilibrium

How does wPNE relate to NE?

Proposition

σ is NE of extensive-form game Γ if and only if there is a belief system μ s.t.

- (i) σ is sequential rational given the system of belief μ at all information sets that are reached given σ ; and
- (ii) the belief system μ is derived through Bayes rule whenever possible given σ .

Note: (i) only require sequential rationality at information sets that are reached (on-path), and

(ii) beliefs at information sets that are reached are correct (coincide with prob. of history being played given σ).

To rule out non-credible threats, we strengthened (i): in wPBE sequential rationality is required at *all* information sets.


Corollary

If (σ, μ) is wPBE of an extensive-form game Γ , then σ is NE of that same game.

wPBE's strategy profile is NE.

Not all NE can be supported (with some belief system) as a wPBE.

Credibility, Part 2



(x, f) is SPNE, but, in a sense, it's a non-credible threat:


Sequential rationality requires P2 to choose a wp1 at $H_2 = \{e_1, e_2\}$ for any belief system μ .

Then, sequential rationality requires P1 to choose e_1 wp1.

Finally, as $P(H_2 \mid (e_1, a)) = 1 > 0$, we have $\mu(H_2)(e_1) = \frac{P(e_1 \mid (e_1, a))}{P(H_2 \mid (e_1, a))} = 1$.

Unique wPBE is $((e_1, a), \mu)$ where $\mu(H_2)(e_1) = 1$.

Credibility, Part 2



If $\mu(H_2)(e_1) \geq 1/2$.

f is sequentially rational $\implies$ P1 chooses x .


$((x, f), \mu)$ is wPBE for $\mu : \mu(H_2)(e_1) \geq 1/2$.

If $\mu(H_2)(e_1) \leq 1/2$.

a is sequentially rational $\implies$ P1 chooses e_1 .

By Bayes' rule, $\mu(H_2)(e_1) = 1 > 1/2$, contradiction!

Credibility, Part 2



If $\mu(H_2)(e_1) = 1/2$.


P2 indifferent between f and a .

P1 chooses x if $0 \geq \max\{\sigma_2(a)3 + (1 - \sigma_2(a))(-1), \sigma_2(a)2 + (1 - \sigma_2(a))(-1)\}$

$$\implies \sigma_2(a) \in [0, 1/4].$$

For any $\sigma_2 : \sigma_2(a) \in [0, 1/4]$, $((x, \sigma_2), \mu)$ is wPBE for $\mu : \mu(H_2)(e_1) = 1/2$.

Credibility, Part 2



If $\mu(H_2)(e_1) = 1/2$.


P2 indifferent between f and a .

P1 never chooses e_2 as if e_2 is preferred to x : $\sigma_2(a)2 + (1 - \sigma_2(a))(-1) \geq 0$,
then e_1 is strictly preferred to both e_2 and x :

$$\sigma_2(a)3 + (1 - \sigma_2(a))(-1) > \sigma_2(a)2 + (1 - \sigma_2(a))(-1) \geq 0$$

If P1 chooses e_1 with positive probability, then by Bayes' rule $\mu(H_2)(e_1) = 1$ a contradiction!

Credibility, Part 2



Conclusion: wPBE are $((\sigma_1, \sigma_2), \mu)$ s.t.

- (i) $\sigma_1(x) = 1$, $\sigma_2(a) = 0$, and $\mu(H_2)(e_1) \in [1/2, 1]$;
- or (ii) $\sigma_1(x) = 1$, $\sigma_2(a) \in [0, 1/4]$, and $\mu(H_2)(e_1) = 1/2$.

Comparing wPBE and SPNE

We already saw a case s.t. σ is SPNE but there is no μ such that (σ, μ) is wPBE.


It is also the case that there wPBE (σ, μ) s.t. σ is not SPNE.

In general:

Remark

Strategy profile that is part of a wPBE need not be an SPNE and a SPNE need not be part of any wPBE.

Comparing wPBE and SPNE



(t, u, r) is unique SPNE but... (b, u, l) can be supported as wPBE.

$p := \mu(\{tu, td\})(tu)$; P3 would choose l over r given p only if

$$p(-1) + (1-p)2 \geq p3 + (1-p)1 \iff 1/5 \geq p.$$

Sequential rationality: P3 chooses l given p ; P2 chooses u (by sequential rationality, never chooses d); P1 chooses b .

So $\forall \mu(\{tu, td\})(tu) \in [0, 1/5]$, $((b, u, l), \mu)$ is wPBE.

Comparing wPBE and SPNE

We already saw a case s.t. σ is SPNE but there is no μ such that (σ, μ) is wPBE.

It is also the case that there wPBE (σ, μ) s.t. σ is not SPNE.

In general:

Remark

Strategy profile that is part of a wPBE need not be an SPNE and a SPNE need be part of any wPBE.

Proposition

In finite extensive-form games of perfect information, set of SPNE is the same as the set of strategy profiles that can be supported as a wPBE (with some belief system).

wPBE strategy profiles and SPNE strategy profiles coincide on games of perfect information, but not necessarily on games of imperfect information

Perfect Bayesian Equilibrium

Why the 'weak' in wPBE? Because wPBE places no restrictions on beliefs in subgames that are not reached given the equilibrium strategy profile.


We here diverge from OR in favor of a more natural definition akin to 'subgame perfection'.

Definition

A strategy profile σ and a belief system μ is a **perfect Bayesian Nash equilibrium** (PBE) (σ, μ) of an extensive-form game Γ if it induces a wPBE in every subgame.

As a wPBE induces a NE, a PBE induces a SPNE.

Not so Perfect PBE



$\forall \mu : \mu(\{he, le\})(he) \in [0, 1], ((x, x), \mu)$ is PBE.

However, any reasonable belief would have $\mu(\{he, le\})(he) = 1/2$.

Problem: can get unreasonable beliefs off-path.

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
5. Applications
6. Beliefs and Sequential Rationality
7. Weak Perfect Bayesian Equilibrium
8. Sequential Equilibrium
 - Sequential Equilibrium
 - Sequential Equilibrium and Trembling-Hand Perfection
9. More

Sequential Equilibrium

Definition

Strategy profile σ and belief system μ is **sequential equilibrium** (SE) (σ, μ) of an extensive-form game Γ if

- (i) σ is sequentially rational given μ ;
- (ii) $\exists$ a sequence of fully mixed strategy profiles $\{\sigma^n\}_n$ inducing a sequence of belief systems μ^n derived through Bayes rule from σ^n s.t. $\sigma^n \rightarrow \sigma$ and $\mu^n \rightarrow \mu$.

Differently from WPBE, SE imposes restrictions on “off-path” beliefs.

Requires beliefs be obtained as a limit of fully mixed beliefs in a way that these beliefs are in the limit consistent with equilibrium play.

Sequential Equilibrium

Theorem

For any finite extensive-form game there is a sequential equilibrium exists.

Proposition

A sequential equilibrium of a finite extensive-form game is also a PBE.

$$(\sigma, \mu) \text{ SE} \implies (\sigma, \mu) \text{ PBE} \implies (\sigma, \mu) \text{ wPBE and } \sigma \text{ SPNE} \implies \sigma \text{ NE}$$

Sequential Equilibrium and Trembling-Hand Perfection

A related notion is that of **extensive-form trembling-hand perfect Nash equilibrium** (ETHPE).

Interpret player choosing at any given information set as a different player.

Define normal-form game of such auxiliary game (the agent normal form of the extensive-form game).

Solve for trembling-hand perfect Nash equilibrium of the auxiliary game.

Sequential Equilibrium and Trembling-Hand Perfection

Proposition

Any extensive-form trembling-hand perfect Nash equilibrium can be supported as a sequential equilibrium by some system of beliefs.

$$\sigma \text{ ETHPE} \implies \exists \mu \text{ s.t. } (\sigma, \mu) \text{ SE}$$

If the extensive-form game is finite, then an ETHPE exists.

Note: A THPE of an extensive-form game need not be subgame perfect.

Overview

1. Why Extensive-Form Games?
2. Extensive-Form Games
3. Nash Equilibrium in Extensive-Form Games
4. Subgame-Perfect Nash Equilibrium
5. Applications
6. Beliefs and Sequential Rationality
7. Weak Perfect Bayesian Equilibrium
8. Sequential Equilibrium
9. More

Off-Path Behaviour and Sequential Rationality

Deviations from sequential equilibria (off-path play) may lead players to question if opponents are sequentially rational.

Reny's (1992 Ecta) Critique of Sequential Equilibrium: SE can rely on “unbelievable” off-path beliefs, as players may hold beliefs that contradict reasonable inferences about past actions.

SE may require beliefs assigning zero prob. to events that seem likely given observed deviations.

Weak Sequential Rationality: Relaxes SE's requirements and tries to address limitations of SE by allowing more plausible off-path beliefs without sacrificing on-path rationality.

Off-Path Behaviour and Sequential Rationality

Deviations from Nash equilibrium (off-path play) may (should?) lead players to question their model of their opponents' behaviour.

Off-path information set could indicate an opponent's mistake.

If an opponent made one mistake, why believe they won't make more?

Unclear how to model this in disciplined manner.

Something to be done here!

The originals: Kreps & Wilson (1982 Ecta); Reny (1992 Ecta).

Strategic Stability and Forward Induction: Fudenberg & Tirole (1991 Book, ch. 11.3); Kohlberg & Mertens (1986 Ecta); Govindan & Wilson (2009 Ecta).

Reputation and Bargaining: Rubinstein (1982 Ecta); Kreps (1982 JET); Abreu & Gul (2000 Ecta); Fudenberg & Tirole (1991 Book, ch. 9).

Limited Foresight: Jehiel & Samet (2007 TE); Ke (2019 TE).

Experiments: McKelvey & Palfrey (1992 Ecta), Brandts, Cabrales & Charness (2008 ET); Cooper & Van Huyck (2003 JET).