Lazy improvement in sequential games

Stéphane Le Roux and Arno Pauly

Université libre de Bruxelles, inVEST project

September 14, 2016

Game theory: for competitive interactions.

Games involve players, strategies, outcomes, preferences.

Game theory: for competitive interactions.

Games involve players, strategies, outcomes, preferences.

Main result in game theory (Nash 1950):
Nash equilibrium (stable strategy combo) in a large class of games.

Game theory: for competitive interactions.

Games involve players, strategies, outcomes, preferences.

Main result in game theory (Nash 1950):
Nash equilibrium (stable strategy combo) in a large class of games.

Problems:

1. Nash equilibrium (NE) requires randomization.

Game theory: for competitive interactions.

Games involve players, strategies, outcomes, preferences.

Main result in game theory (Nash 1950):
Nash equilibrium (stable strategy combo) in a large class of games.

Problems:

1. Nash equilibrium (NE) requires randomization.
2. Why would players play NE?

Game theory: for competitive interactions.

Games involve players, strategies, outcomes, preferences.

Main result in game theory (Nash 1950):
Nash equilibrium (stable strategy combo) in a large class of games.

Problems:

1. Nash equilibrium (NE) requires randomization.
2. Why would players play NE?

Partial solution (Kukushkin 2002):
Finite convergence to pure NE in finite sequential games.

Game theory: for competitive interactions.

Games involve players, strategies, outcomes, preferences.

Main result in game theory (Nash 1950):
Nash equilibrium (stable strategy combo) in a large class of games.

Problems:

1. Nash equilibrium (NE) requires randomization.
2. Why would players play NE?

Partial solution (Kukushkin 2002):
Finite convergence to pure NE in finite sequential games.

We generalize 2002 for finite and infinite sequential games.

Finite sequential games

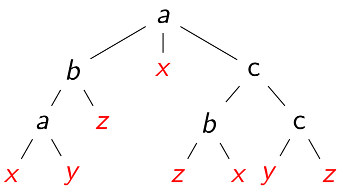

a,b,c are players

x, y, z are outcomes.

A preference is a binary relation over the outcomes.

Finite sequential games

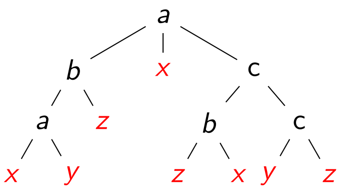

a,b,c are players

x, y, z are outcomes.

A preference is a binary relation over the outcomes.

Classical outcome $(2, 7, 4)$: player a gets 2, player b gets 7...

Finite sequential games

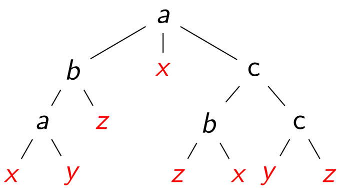

a,b,c are players

x, y, z are outcomes.

A preference is a binary relation over the outcomes.

Classical outcome $(2, 7, 4)$: player a gets 2, player b gets 7...

Classical preference: $(0, 2, 1) \prec_b (9, 3, 0)$.

Strategies and strategy profiles

Double lines represent strategical choices.

Strategy for player a

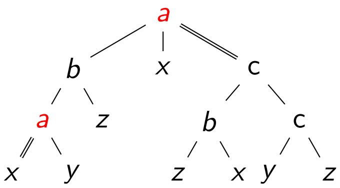

Strategies and strategy profiles

Double lines represent strategical choices.

Strategy for player *a*

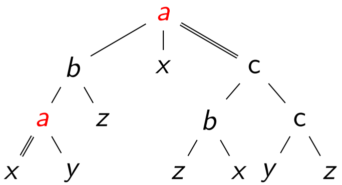

Strategy profile

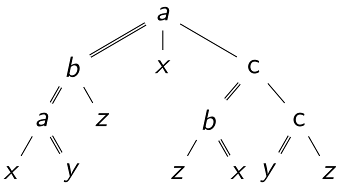

Strategies and strategy profiles

Double lines represent strategical choices.

Strategy for player a

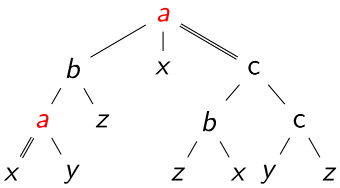

Strategy profile, induced outcome

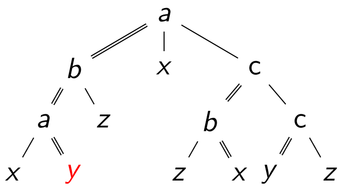

Nash equilibrium

Definition by example

Game

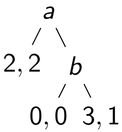

¬NE

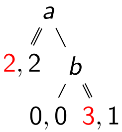

Nash equilibrium

Definition by example

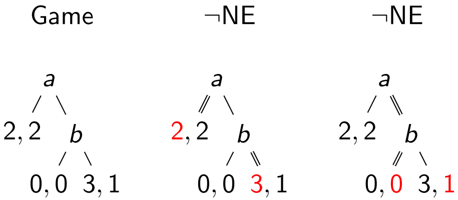

Nash equilibrium

Definition by example

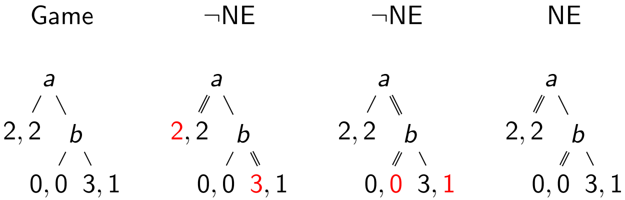

Nash equilibrium

Definition by example

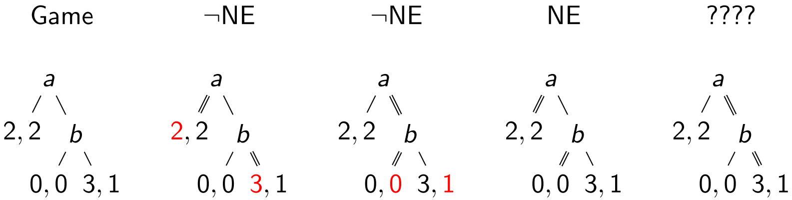

Nash equilibrium

Definition by example

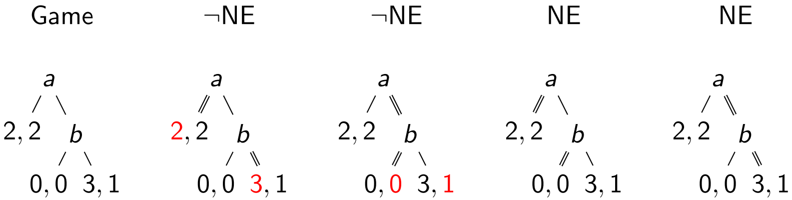

Nash equilibrium

Definition by example

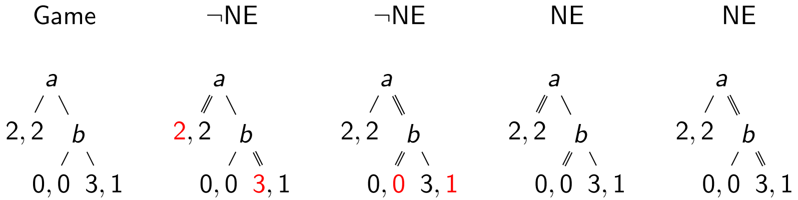

1953 Kuhn: all finite sequential games have NE.

1995 Aumann: in some games, some NE correspond to common knowledge of rationality.

Nash equilibrium

Definition by example

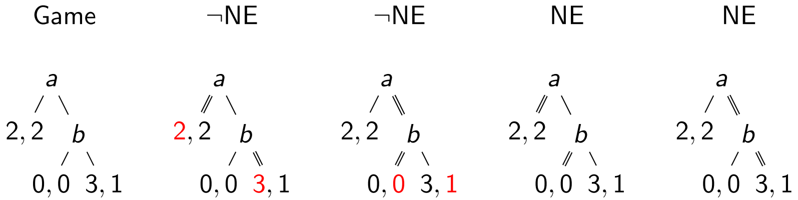

1953 Kuhn: all finite sequential games have NE.

1995 Aumann: in some games, some NE correspond to common knowledge of rationality.

- How about the other games?

Nash equilibrium

Definition by example

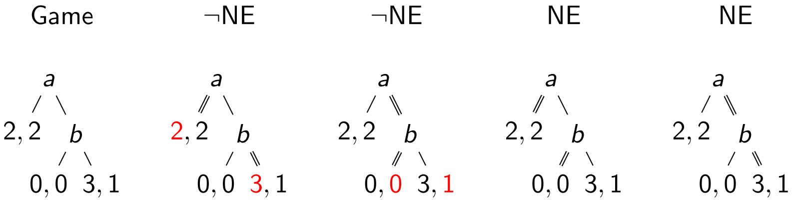

1953 Kuhn: all finite sequential games have NE.

1995 Aumann: in some games, some NE correspond to common knowledge of rationality.

- How about the other games?
- What if some players are not knowledgeable?

Nash equilibrium

Definition by example

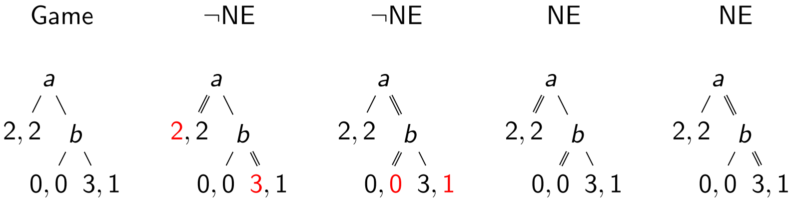

1953 Kuhn: all finite sequential games have NE.

1995 Aumann: in some games, some NE correspond to common knowledge of rationality.

- How about the other games?
- What if some players are not knowledgeable?
- What if some players are not rational?

Nash Equilibrium and better-response dynamics

- For all non-NE, the better-response lets one player improve it.

Nash Equilibrium and better-response dynamics

- For all non-NE, the better-response lets one player improve it.
- NE(s) $\Leftrightarrow$ s is a sink for the better-response dynamics.

Nash Equilibrium and better-response dynamics

- For all non-NE, the better-response lets one player improve it.
- NE(s) $\Leftrightarrow$ s is a sink for the better-response dynamics.

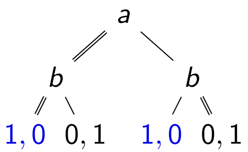

Nash Equilibrium and better-response dynamics

- For all non-NE, the better-response lets one player improve it.
- NE(s) $\Leftrightarrow$ s is a sink for the better-response dynamics.

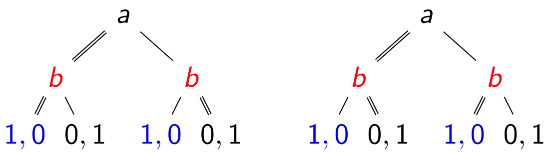

Nash Equilibrium and better-response dynamics

- For all non-NE, the better-response lets one player improve it.
- NE(s) $\Leftrightarrow$ s is a sink for the better-response dynamics.

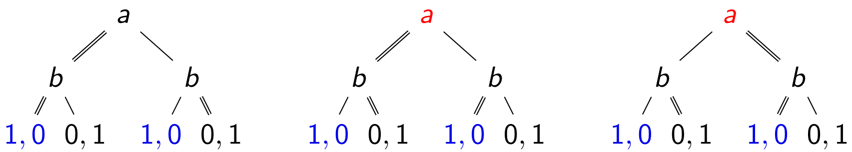

Nash Equilibrium and better-response dynamics

- For all non-NE, the better-response lets one player improve it.
- NE(s) ⇔ s is a sink for the better-response dynamics.

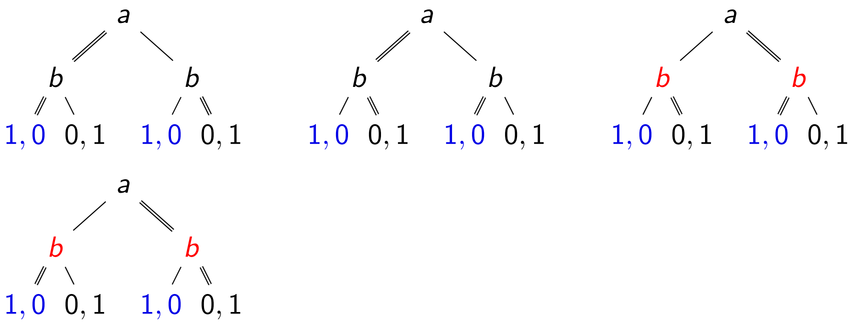

Nash Equilibrium and better-response dynamics

- For all non-NE, the better-response lets one player improve it.
- NE(s) $\Leftrightarrow$ s is a sink for the better-response dynamics.

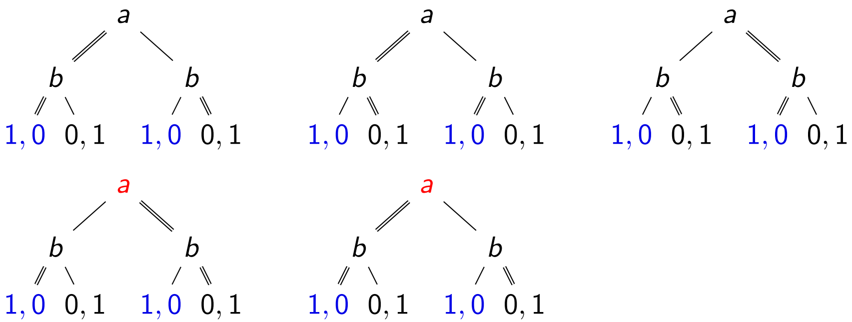

Lazy convertibility

Restricting how players change strategies

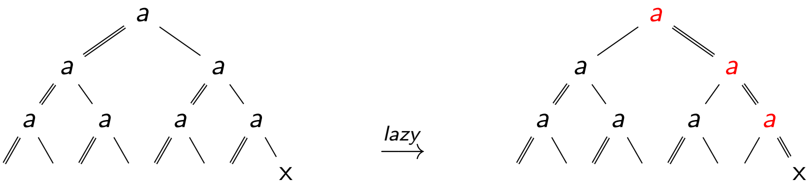

Lazy convertibility

Restricting how players change strategies

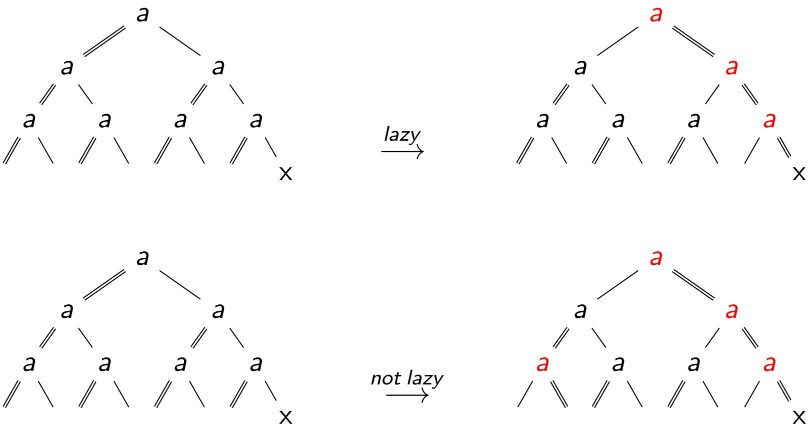

Lazy convertibility

Restricting how players change strategies

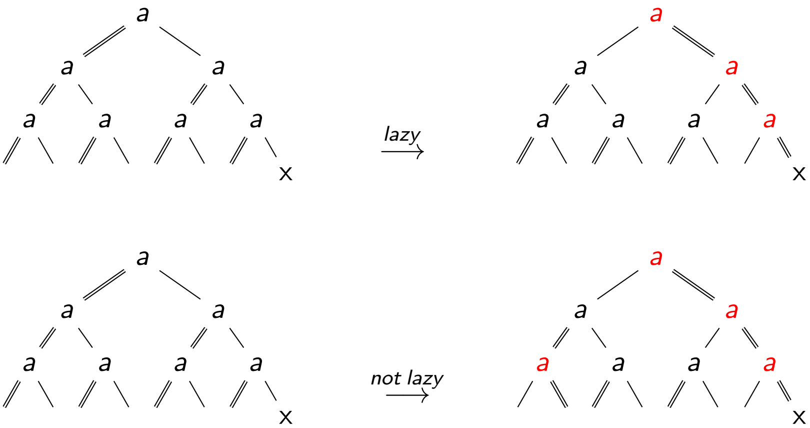

If a player can convert and reach a leaf, she can do so lazily.

Lazy improvement (lazy better-response dynamics)

For all non-NE, let a player improve by lazy convertibility.

Lazy improvement (lazy better-response dynamics)

For all non-NE, let a player improve by lazy convertibility.

The NE are the sinks of lazy improvement (LI).

Lazy improvement (lazy better-response dynamics)

For all non-NE, let a player improve by lazy convertibility.

The NE are the sinks of lazy improvement (LI).

$\Rightarrow$ $LI \subseteq BR$ so every NE is a sink of LI.

Lazy improvement (lazy better-response dynamics)

For all non-NE, let a player improve by lazy convertibility.

The NE are the sinks of lazy improvement (LI).

$\Rightarrow$ $LI \subseteq BR$ so every NE is a sink of LI.

$\Leftarrow$ If a player can improve, she can do it lazily.

Lazy improvement (lazy better-response dynamics)

For all non-NE, let a player improve by lazy convertibility.

The NE are the sinks of lazy improvement (LI).

$\Rightarrow$ $LI \subseteq BR$ so every NE is a sink of LI.

$\Leftarrow$ If a player can improve, she can do it lazily.

Theorem (Kukushkin, 2002)

For games with $\mathbb{R}$-valued payoffs, LI terminates.

Lazy improvement (lazy better-response dynamics)

For all non-NE, let a player improve by lazy convertibility.

The NE are the sinks of lazy improvement (LI).

$\Rightarrow$ $LI \subseteq BR$ so every NE is a sink of LI.

$\Leftarrow$ If a player can improve, she can do it lazily.

Theorem (Kukushkin, 2002)

For games with $\mathbb{R}$-valued payoffs, LI terminates.

Theorem (New)

Consider a game with n nodes.

- *If no player has a preference path longer than h, LI terminates within $h \cdot n$ steps.*

Lazy improvement (lazy better-response dynamics)

For all non-NE, let a player improve by lazy convertibility.

The NE are the sinks of lazy improvement (LI).

$\Rightarrow$ $LI \subseteq BR$ so every NE is a sink of LI.

$\Leftarrow$ If a player can improve, she can do it lazily.

Theorem (Kukushkin, 2002)

For games with $\mathbb{R}$-valued payoffs, LI terminates.

Theorem (New)

Consider a game with n nodes.

- *If no player has a preference path longer than h, LI terminates within $h \cdot n$ steps.*
- *If a player has no preference path longer than h, she makes less than $h \cdot n$ steps in any LI sequence.*

Idea of the proof

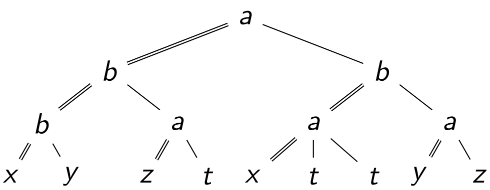

Idea of the proof

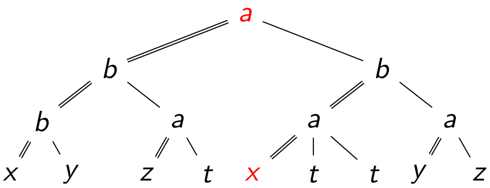

Which player avoids which outcomes (multiset):

	x	y	z	t
a	1	0	1	3
b	0	2	1	0

Idea of the proof

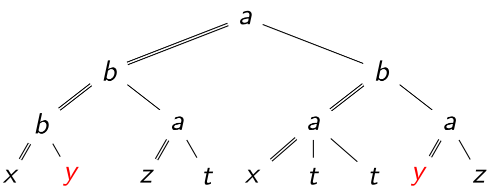

Which player avoids which outcomes (multiset):

	x	y	z	t
a	1	0	1	3
b	0	2	1	0

Idea of the proof

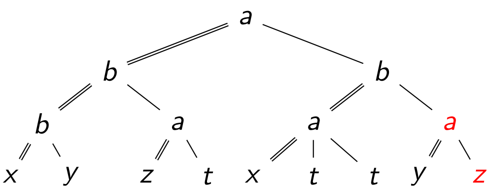

Which player avoids which outcomes (multiset):

	x	y	z	t
a	1	0	1	3
b	0	2	1	0

Idea of the proof

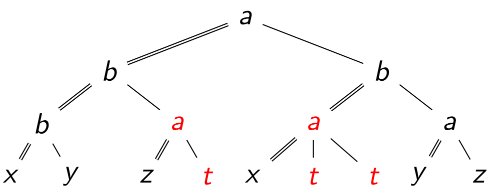

Which player avoids which outcomes (multiset):

	x	y	z	t
a	1	0	1	3
b	0	2	1	0

Idea of the proof

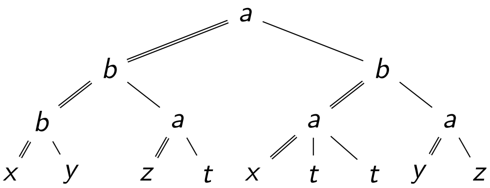

Which player avoids which outcomes (multiset):

	x	y	z	t
a	1	0	1	3
b	0	2	1	0

- Lazy convertibility by Player *a* doesn't affect *b*'s multiset.

Idea of the proof

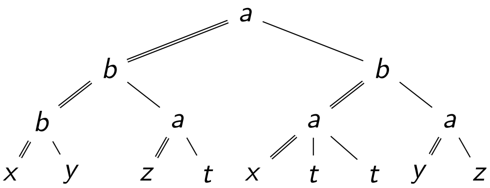

Which player avoids which outcomes (multiset):

	x	y	z	t
a	1	0	1	3
b	0	2	1	0

- Lazy convertibility by Player *a* doesn't affect *b*'s multiset.
- LI by a player decreases her multiset (wrt her preference).

Idea of the proof

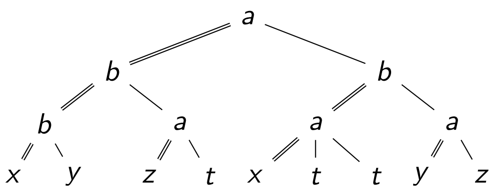

Which player avoids which outcomes (multiset):

	x	y	z	t
a	1	0	1	3
b	0	2	1	0

- Lazy convertibility by Player *a* doesn't affect *b*'s multiset.
- LI by a player decreases her multiset (wrt her preference).
- $\Rightarrow$ self-stabilzation of players with acyclic preference.

Some properties of lazy improvement

1. No probability.

Some properties of lazy improvement

1. No probability.
2. Acyclic preferences (rationality) suffices for termination.
3. Quick stabilisation.
4. Rational players stabilize regardless of opponents.

Some properties of lazy improvement

1. No probability.
2. Acyclic preferences (rationality) suffices for termination.
3. Quick stabilisation.
4. Rational players stabilize regardless of opponents.
5. LI step independent of previous LI steps.
6. A players needn't know she has opponents.
7. No need to know all ones options: maximizing not required.

Some properties of lazy improvement

1. No probability.
2. Acyclic preferences (rationality) suffices for termination.
3. Quick stabilisation.
4. Rational players stabilize regardless of opponents.
5. LI step independent of previous LI steps.
6. A players needn't know she has opponents.
7. No need to know all ones options: maximizing not required.
8. A new proof technique for existence of NE.

Lazy improvement in infinite sequential games

Infinite sequential games

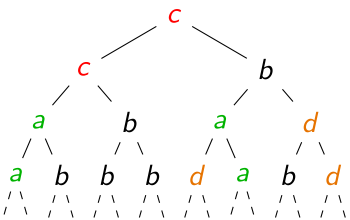

- No leaf: outcomes are attached to infinite paths.
- Strategy and NE are defined as in the finite case.

Infinite sequential games

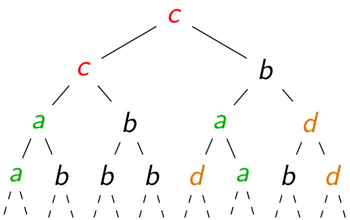

- No leaf: outcomes are attached to infinite paths.
- Strategy and NE are defined as in the finite case.

Proposition

If the player preferences are all acyclic, so is lazy improvement.

Lazy improvement has ω-cycles

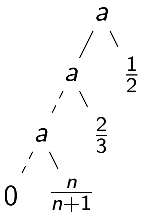

Lazy improvement has ω-cycles

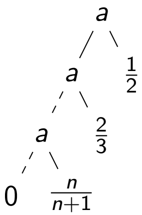

The limit of the sequence is its first element.

Lazy improvement has ω-cycles

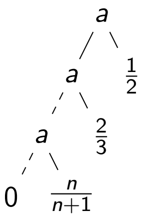

The limit of the sequence is its first element.

In the remainder, the game has finite branching, finitely many players, and continuous real-valued payoffs.

Lazy improvement has ω-cycles

The limit of the sequence is its first element.

In the remainder, the game has finite branching, finitely many players, and continuous real-valued payoffs.

Lazy improvement has ω-cycles

The limit of the sequence is its first element.

In the remainder, the game has finite branching, finitely many players, and continuous real-valued payoffs.

Theorem

Lazy ϵ-improvement terminates on ϵ-NE.

Deepening lazy improvement

Algorithm:

1. $n := 1$;
2. Let s be any strategy profile;
3. Repeat {
4. Let the players improve s by lazy convertibility above depth n, until some stable s_n;
5. $s := s_n$;
6. $n := n + 1$; }

Theorem

All accumulation points of $(s_n)_{n \in \mathbb{N}}$ are NE.

Conclusion

Existence of randomized NE (1950) raised two issues: randomization and convergence.

Kukushkin (2002) addressed these issues partially: LI in finite sequential games with $\mathbb{R}$-valued payoffs terminates on NE.

Conclusion

Existence of randomized NE (1950) raised two issues: randomization and convergence.

Kukushkin (2002) addressed these issues partially: LI in finite sequential games with $\mathbb{R}$-valued payoffs terminates on NE.

We contributed on

- Finite games: termination bounds + self-stablization.
- Infinite games: several directions, assuming continuity, finite branching and number of players.

Fair lazy improvement

A sequence of improvements is fair if the following holds: for all positive reals r, if improvements by more than r are possible infinitely often during the sequence, they also occur infinitely often.

Proposition

Fair lazy improvement

A sequence of improvements is fair if the following holds: for all positive reals r, if improvements by more than r are possible infinitely often during the sequence, they also occur infinitely often.

Proposition

- *One accumulation point of fair lazy improvement is NE.*

Fair lazy improvement

A sequence of improvements is fair if the following holds: for all positive reals r, if improvements by more than r are possible infinitely often during the sequence, they also occur infinitely often.

Proposition

- *One accumulation point of fair lazy improvement is NE.*
- *If the accumulation points are finitely many, there are all NE.*

Fair lazy improvement

A sequence of improvements is fair if the following holds: for all positive reals r, if improvements by more than r are possible infinitely often during the sequence, they also occur infinitely often.

Proposition

- *One accumulation point of fair lazy improvement is NE.*
- *If the accumulation points are finitely many, there are all NE.*
- *If the payoffs are (suitably) Lipschitz, the points are all NE.*

Fair lazy improvement

A sequence of improvements is fair if the following holds: for all positive reals r, if improvements by more than r are possible infinitely often during the sequence, they also occur infinitely often.

Proposition

- *One accumulation point of fair lazy improvement is NE.*
- *If the accumulation points are finitely many, there are all NE.*
- *If the payoffs are (suitably) Lipschitz, the points are all NE.*

Open question: are all accumulation points of fair LI NE?

Ordinal lazy improvement

Algorithm:

1. Let s be any strategy profile;
2. Repeat $\{$
3. Let the players improve s lazily;
4. Let s be an accumulation point of the sequence;$\}$

A Δ_2^0 set is a countable union of closed sets and is also a countable intersection of open sets.

Proposition

If the players have Δ_2^0 boolean objectives, the sequence reaches an NE after a countable ordinal of steps.